National Aeronautics and
Space Administration

HIGH-END COMPUTING CAPABILITY PORTFOLIO

William Thigpen
NASA Advanced Supercomputing Division

April 10, 2021

New Endeavour Supercomputer Now in Production

- A new HPE Superdome Flex system with 32 Cascade Lake processors was released into production on March 31, 2021.
  - Endeavour3 (released in production) and Endeavour4 (not released) will have a total of 1,792 cores and 12 TB of memory, with a Standard Billing Unit (SBU) value of 1.31.
  - The original Endeavour systems were decommissioned when the new system was released into production.
- The original Endeavour systems were palced in production with Sandy Bridge processors in February 2013.
  - Endeavour1 and Endeavour2 had a total of 1,536 cores with 6 TB of memory.
  - The shared memory Endeavour system was installed to take the place of the Columbia supercomputer, which was decommissioned in 2013.
  - When first installed, these systems had more processing power than what remained of Columbia but using ten percent of the physical space.
- The global shared memory architecture allows each processor to access the memory of the other processors in the system, enabling users to run large-memory or data-intensive jobs that cannot be run on other HECC systems.

** 1 SBU equals 1 hour of a Pleiades Broadwell 28-core node.*

IMPACT: The replacement systems increase Endeavour's compute capabilities by 483%, enabling quicker turnaround time for HECC scientists with large, shared memory capacity jobs.

HPE Superdome Flex eight-chassis rack. *Image courtesy of HPE*

HECC Awards Allocations for AWS Cloud Resources

- The HECC Cloud team conducted a Call For Proposals looking for innovative work to use cloud-based resources for science and engineering applications.
- Five projects were selected to use HECC Cloud resources at Amazon Web Services (AWS) with HECC funds:
  - Cloud Compute for the Solar Dynamics Observatory: PI, Mark Cheung (Lockheed Martin)
  - Prototype Science Data Analysis Pipeline on HECC Cloud; PI, Jon Jenkins (Ames)
  - Cloud Computing Replacement for Render Farm; PI, Paul Kessler (Langley).
  - AWS Implementation of Methane Point Source Workflow: PI, Andrew Thorpe (JPL)
  - TRopospheric Ozone and Precursors from Earth System Sounding (TROPESS): PI. Adrian Tinio (JPL)
- The Cloud team has begun work with the PIs to set up AWS accounts and configure resources appropriately for each project.

IMPACT: Access to compute resources in the cloud can improve user efficiency by reducing wait times. It can also be more efficient, such as in use cases where data is cloud-resident, by eliminating data transfers.

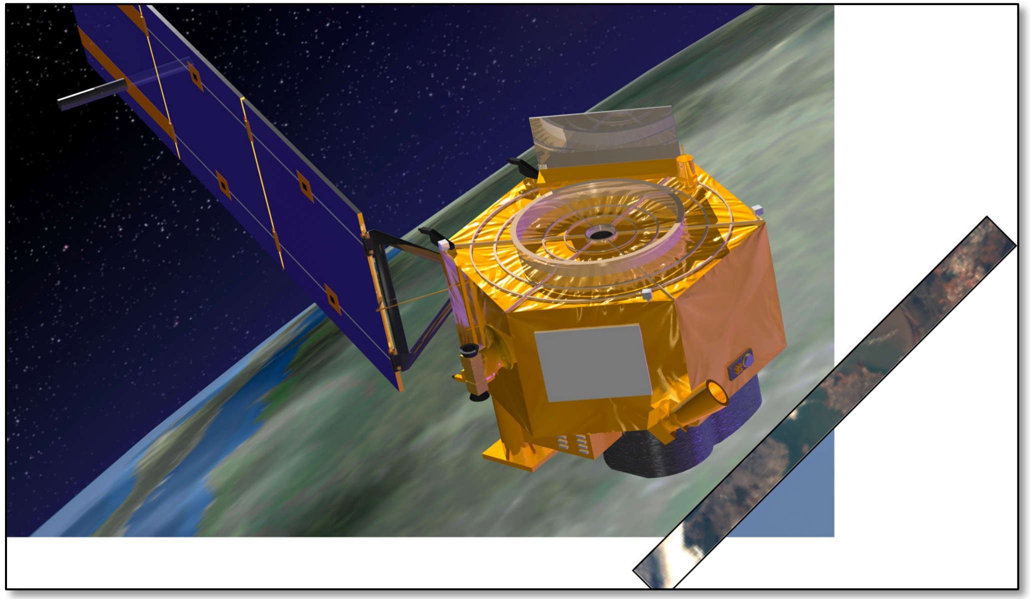

Artist's conception of the EO-1 satellite, inset with an image of San Pablo Bay, CA from its Hyperion instrument, which collected 220 unique spectral channels 0.357–2.576 µm with a 10-nm bandwidth. One cloud project (PI Jon Jenkins) will port the science data pipeline for Hyperion to AWS to help inform decisions about the upcoming Surface Biology and Geology science mission. *NASA*

New Secure Front-End Systems are Operational

- The HECC Systems team upgraded the Secure Front-End (SFE) systems, increasing security and maintainability. Work included:
  - Migrated the SFEs to the NASA-approved RedHat Linux.
  - Removed the licensed Tectia SSH server in favor of the now supported OpenSSH server. This improvement eliminates the use of obsolete SSH algorithms; supports SecurID together with public key, password, or personal identity verification (PIV) authentication in a single SSH server instead of two; and eliminates licensing costs of the Tectia SSH server.
- The systems are now current with latest federal Cybersecurity Standards and Engineering Team (CSET) mandates.
- The new SFEs are implemented as four virtual machines running on two physical hypervisor systems to simplify system management.
- The systems are now based on the existing Secure Unattended Proxy (SUP) architecture, which provides a unified architecture across all bastions to reduce administrative burden, and fully constrains users to exactly the operations needed to provide SFE functionality.

IMPACT: The new Secure Front-End systems provide increased security for the HECC environment through up-to-date SSH algorithms, better adherence to NASA security requirements, and a more easily supported operating system.

The SFEs provide a critical piece of the current NAS HECC security architecture providing a secure bastion between the NAS internal network and the internet.

HECC Supercomputer Usage Sets New Record*

- In March 2021, the combined usage of HECC supercomputers set a new record of 11,662,724 Standard Billing Units (SBUs).*
- The usage by 416 of NASA's science and engineering groups exceeded the previous record of 9,677,882 SBUs set in April 2020 by 1,984,842 SBUs (20%).
- The record was achieved in great part by the Science Mission Directorate's Earth Science group from a coupled Goddard Earth Observing System (GEOS)/Estimating the Circulation and Climate of the Ocean (ECCO).
- Usage of Pleaides, Aitken, Electra, Merope, and Endeavour contributed to this record. The new record was enabled by the addition of 1,024 Aitken Rome nodes, with 128 cores per node.
- The top 10 projects' usage ranged between 214,515 and 712,752 SBUs, and together accounted for over 33% of the total usage.
- The HECC Project continues to evaluate and plan resources to address the future requirements of NASA's users.

IMPACT: Both the increased capacity of HECC systems and support work with users to optimize their run capacities provides mission directorates with more resources to accomplish their goals and objectives.

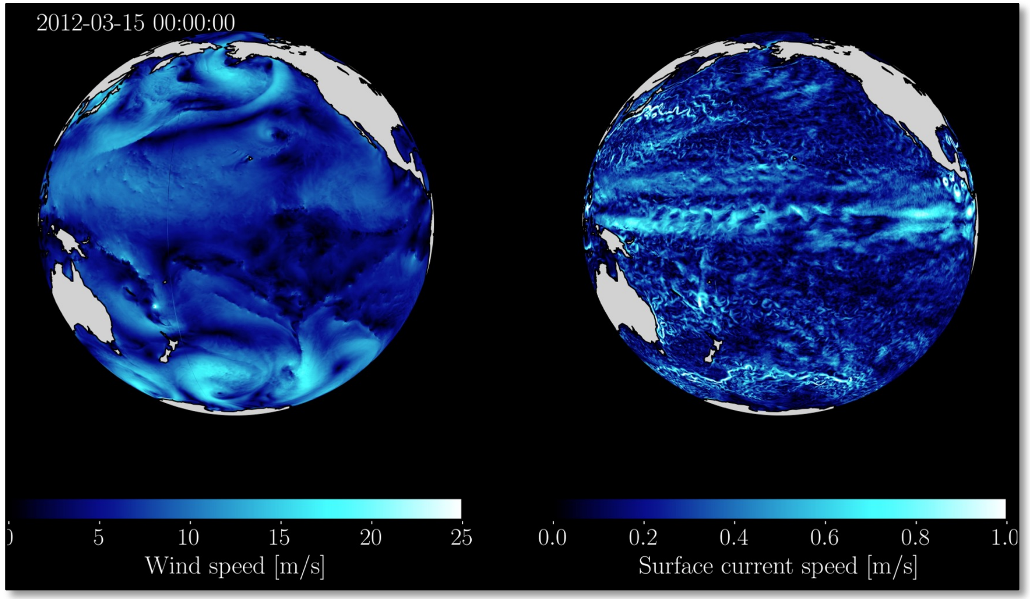

Snapshots of near-surface wind speed (left panel) and ocean-surface-current speed (right panel) from a coupled Goddard Earth Observing System (GEOS)/Estimating the Circulation and Climate of the Ocean (ECCO) numerical simulation. *Ehud Strobach, University of Maryland; Andrea Molod, NASA/Goddard*

** 1 SBU represents the work that can be done in 1 hour on a Pleiades Broadwell 28-core node.*

Vis Team Expands Access to 1/48° ECCO Ocean Simulation

- The HECC Visualization and Analysis and Big Data teams deployed a new database of visualizations from the 1/48° ECCO project's 14-month ocean simulation.
  - A total of 720,235 animations are available, covering all of the fields and ocean depths for the entire ocean, except for the Arctic region, and at different resolutions and speeds (different time steps).
  - A web page with interactive Javascript and image thumbnails make it simple to find an animation: https://data.nas.nasa.gov/viz/vizdata/llc4320/index.html
- Previously, visualizing the 1/48° ECCO data was difficult, as it required a software installation, and visualizations over time might require several terabytes of downloads.
- In the first two months of availability the team collected these statistics about animation downloads:
  - 63 unique IP addresses, 2000+ hits, 320 GBs downloaded.
  - The peak day had 500 hits and 130 GBs downloaded.
- The ECCO group created a web page describing in detail how to access the database: https://ecco-group.org/world-of-ecco.htm
- Adding Arctic visualizations to the database is in progress.

IMPACT: The new database created by HECC visualization experts allows a broader community of researchers to explore time-varying 1/48° ECCO simulation data, as the precomputed visualizations require much less effort to view.

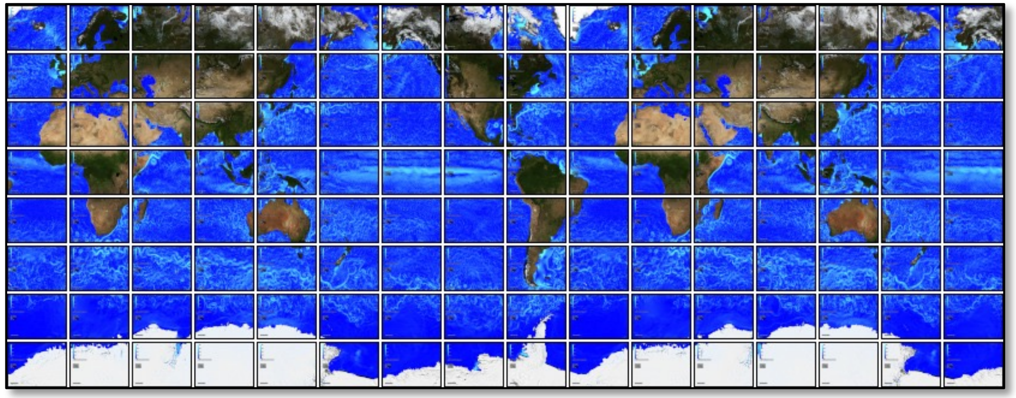

An image from the HECC-developed web page showing the available animations of the ocean speed at the surface from the Estimating the Circulation and Climate of the Ocean (ECCO) project's simulation output. Each small image in the grid can be clicked to see a range of visualizations at different resolutions and speeds for that location. *David Ellsworth, NASA/Ames*

Identifying the Genetic Ancestry of COVID-19 Patients from New York City*

- Scientists on the COVID-19 International Research Team (COV-IRT), co-led by NASA Ames and Rice University, performed RNA-sequencing (RNA-Seq) experiments for samples taken from 732 COVID-19 patients in New York City.
- The team applied machine learning to genetic ancestry principal components in order to classify COV-IRT patients into five continental ancestry groups: Admixed American, African, East Asian, European, and South Asian. African ancestry individuals could be further distinguished as Hispanic or non-Hispanic, based on their levels of Native American ancestry.
- Findings allow members of COV-IRT to further assess how COVID-19 impacts patients based on their genetic ancestry, which could lead to more effective treatments and means of prevention.
- HECC resources and services were critical in generating the data for this project and to develop and test three data processing pipelines. HECC experts were instrumental in helping troubleshoot, debug, and optimize the PBS scripts within each pipeline, allowing for parallel processing of samples and reduced processing time.

** HECC provided supercomputing resources and services in support of this work*

IMPACT: As part of NASA's response to the COVID-19 pandemic, HECC provides supercomputing resources and expertise to support researchers in the fight to understand the SARS-CoV-2 virus, and to develop treatments and vaccines.

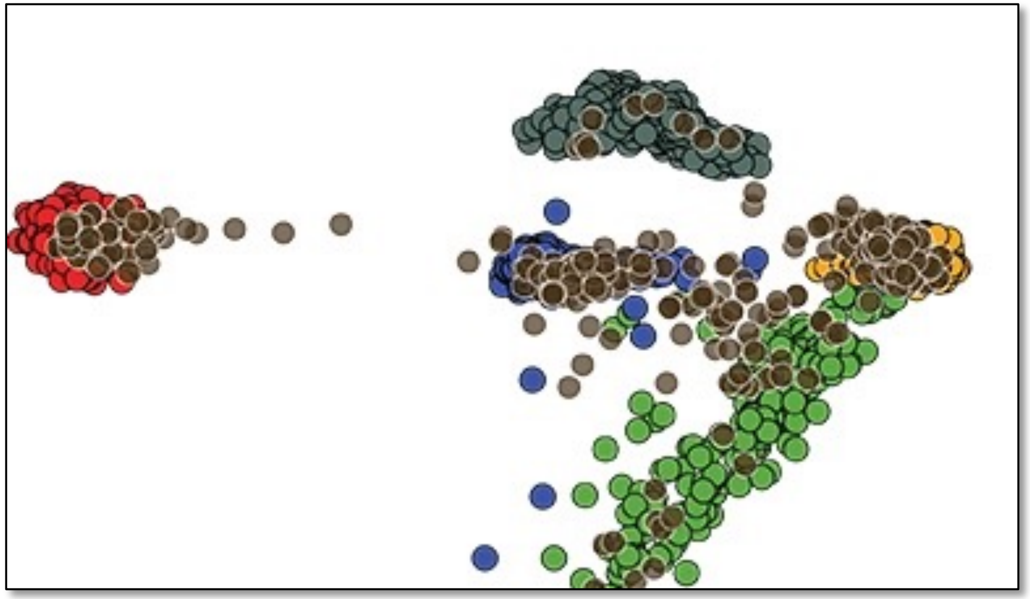

Animation showing an interactive PCA plot, which places COVID-19 International Research Team (COV-IRT) patients relative to reference populations from Africa, the Americas, East Asia, Europe, and South Asia. *Shashwat Deepali Nagar, Georgia Institute of Technology; Andrew B. Conley, Applied Bioinformatics Laboratory*

Predicting Jet Noise for Full-Scale Low-Boom Aircraft*

- Researchers at NASA Ames ran computational aeroacoustic simulations on Pleiades and Electra, using their Launch Ascent and Vehicle Aerodynamics (LAVA) solver, to improve the solver's wall-modeled large eddy simulation (WMLES) capability—a step toward predicting jet acoustics for NASA's Low Boom Flight Simulator (LBFD).
- In addition to reducing sonic boom, a supersonic airplane must also satisfy noise constraints while traveling at subsonic speeds (during takeoff and landing) in order to be certified. Computational aeroacoustic simulation tools can be used to assess new designs for the LBFD at these lower speeds.
  - Most computational fluid dynamics (CFD) simulations rely on Reynolds-Averaged Navier-Stokes (RANS) methods and are still limited to simple geometries; the scale-resolving simulations needed to predict accurate far-field noise are too computationally demanding to investigate full airframe interaction noise.
  - To address this, the LAVA team explored a new strategy to achieve full-scale airframe simulations that include jet-noise prediction with complex geometries like multi-stream chevron nozzles by using WMLES.
  - The researchers ran scale-resolving simulations with LAVA, scrutinizing hybrid RANS/LES methods on validation cases with increasing complexity, including jet surface interaction noise, multi-stream, and chevron nozzle designs.
- The team's simulations successfully predicted jet noise as well as jet surface interaction noise. Findings have already resulted in new insights and additional experimental investigations of certain flow phenomena.

IMPACT: Gaining confidence in the predictive capabilities of hybrid Reynolds-Averaged Navier-Stokes (RANS)-large eddy simulation (LES) for jet noise is a step toward vehicle certification through computational aeroacoustic analysis. This work supports NASA's Commercial Supersonic Transport Project.

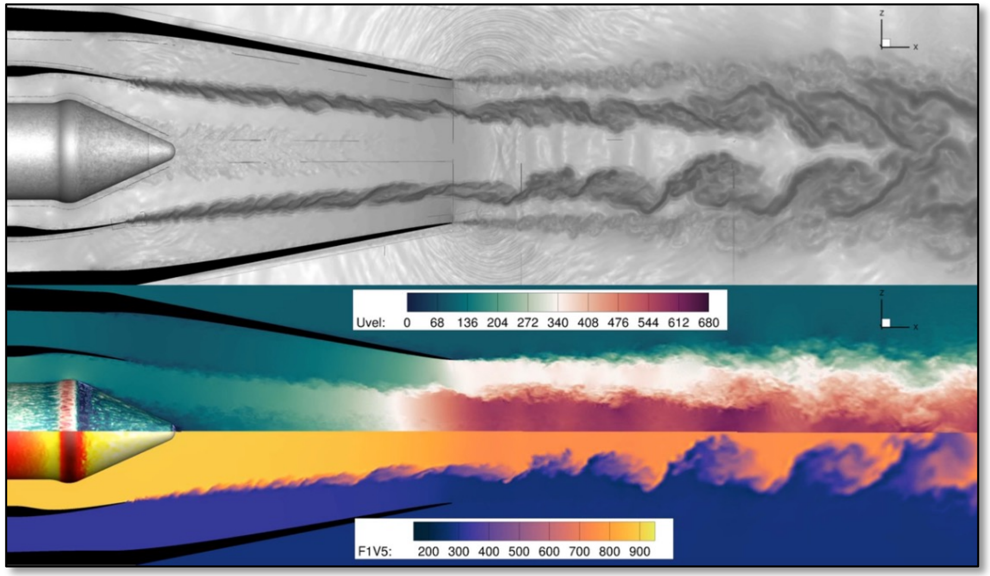

WMLES simulation of two-stream heated internal-plug nozzle design performed with the LAVA solver. Top: Density gradient magnitude. Middle: Stream-wise velocity. Bottom: Temperature. *Gerrit-Daniel Stich, Timothy Sandstrom, NASA/Ames*

** HECC provided supercomputing resources and services in support of this work*

Papers

- “**A Transiting Warm Giant Planet Around the Young Active Star TOI-201**,” M. Hobson, et al., arXiv:2103.02685 [astro-ph.EP], March 3, 2021. *
  https://arxiv.org/abs/2103.02685
- “**Analysis of Droplet Evaporation in Isotropic Turbulence Through Droplet-Resolved DNS**,” M. Dodd, D. Mohaddes, A. Ferrante, M. Ihme, International Journal of Heat and Mass Transfer, vol. 172, published online March 7, 2021. *
  https://www.sciencedirect.com/science/article/abs/pii/S001793102100260X
- “**Effects of the Nanotube Length and Network Morphology on the Deformation Mechanisms and Mechanical Properties of Cross-Linked Carbon Nanotube Films**,” A. H. Banna, et al., Journal of Applied Physis, vol. 129, published online March 8, 2021. *
  https://aip.scitation.org/doi/abs/10.1063/5.0033442
- “**Magnetotail-Inner Magnetosphere Transport Associated with Fast Flows Based on Combined Global-Hybrid and CIMI Simulation**,” Y. Lin, et al., Journal of Geophysical Research: Space Physics, vol. 126, issue 3, March 8, 2021. *
  https://agupubs.onlinelibrary.wiley.com/doi/full/10.1029/2020JA028405
- “**Efficient Projection Kernels for Discontinuous Galerkin Simulations of Disperse Multiphase Flows on Arbitrary Curved Elements**,” E. Ching, M. Ihme, Journal of Computational Physics, vol. 435, March 9, 2021. *
  https://www.sciencedirect.com/science/article/pii/S0021999121001613
- “**Configuration of the Earth’s Magnetotail Current Sheet**,” A. Aremyev, et al., Geophysical Research Letters, vol. 48, issue 6, March 10, 2021. *
  https://agupubs.onlinelibrary.wiley.com/doi/full/10.1029/2020GL092153

** HECC provided supercomputing resources and services in support of this work*

Papers (cont.)

- "**Magnetopause Reconnection and Indents Caused by Foreshock Turbulence**," L.-J. Chen, J. Ng, Y. Omelchenko, S. Wang, arXiv:2103.07448 [physics.space-ph], March 12, 2021. *
  https://arxiv.org/abs/2103.07448
- "**Computational Electrophysiology from a Single Molecular Dynamics Simulation and the Electrodiffusion Model**," A. Pohorille, M. Wilson, Journal of Physical Chemistry B, vol. 125, published online March 17, 2021. *
  https://pubs.acs.org/doi/abs/10.1021/acs.jpcb.0c10737
- "**Two Bright M Dwarfs Hosting Ultra-Short-Period Super-Earths with Earth-Like Compositions**," T. Hirano, et al., arXiv:2103.12760 [astro-ph.EP], March 23, 2021. *
  https://arxiv.org/abs/2103.12760
- "**TOI-1634 b: An Ultra-Short Period Keystone Planet Sitting Inside the M Dwarf Radius Valley**," R. Cloutier, et al., arXiv:2103.12790, [astro-ph.EP], March 23, 2021. *
  https://arxiv.org/abs/2103.12790
- "**The TESS Objects of Interest Catalog from the TESS Prime Mission**," N. Guerrero, et al., arXiv:2103.12538 [astro-ph.EP], March 23, 2021. *
  https://arxiv.org/abs/2103.12538
- "**Sensitivity of Hypersonic Dusty Flows to Physical Modeling of the Particle Phase**," E. Ching, M. Barnhardt, M. Ihme, Journal of Spacecraft and Rockets, published online March 25, 2021. *
  https://arc.aiaa.org/doi/full/10.2514/1.A34810

** HECC provided supercomputing resources and services in support of this work*

Papers (cont.)

- "**A Discontinuous Galerkin Method for Wall-Modeled Large Eddy Simulations**," Y. Lv, et al., Computers & Fluids, vol. 222, published online March 25, 2021. *
  https://www.sciencedirect.com/science/article/abs/pii/S0045793021000992
- "**Development of a Particle Collision Algorithm for Discontinuous Galerkin Simulations of Compressible Multiphase Flows**," E. Ching, M. Ihme, Journal of Computational Physics, vol. 436, published online March 26, 2021. *
  https://www.sciencedirect.com/science/article/pii/S002199912100214X
- "**Origins and Demographics of Wandering Black Holes**," A. Ricarte, et al., Monthly Notices of the Royal Astronomical Society, published online March 26, 2021. *
  https://academic.oup.com/mnras/advance-article/doi/10.1093/mnras/stab866/6189716
- "**TESS Re-Observes the Young Multi-Planet System TOI-451: Refined Ephemeris and Activity Evolution**," O. Barragán, et al., Research Notes of the American Astronomical Society, vol. 5, no. 3, March 2021. *
  https://iopscience.iop.org/article/10.3847/2515-5172/abef70/meta

** HECC provided supercomputing resources and services in support of this work*

Presentations

- "**Fine Pointing of Laser Beams by Using Laser Arrays for Applications to CubeSats**," P. Goorjian, presented at the 2021 SPIE LASE, published online in Proceedings Vol. 11678, Free-Space Laser Communications XXXIII, March 5, 2021. *https://www.spiedigitallibrary.org/conference-proceedings-of-spie/11678/116780E/Fine-pointing-of-laser-beams-by-using-laser-arrays-for/10.1117/12.2575661.full?SSO=1*

** HECC provided supercomputing resources and services in support of this work*

News and Events

- **Deciphering Medical Risk with High-Performance Computing**, *Medical Design Briefs*, March 1, 2021—Researchers at NASA Ames are performing DNA sequence analysis on the Pleiades supercomputer to help identify medical risks to astronauts during future missions to the Moon, and to answer questions about COVID-19 genetic variations in humans. *https://www.medicaldesignbriefs.com/component/content/article/mdb/tech-briefs/38751*
- **Using NAS-Developed Tools to Quiet the Boom of Supersonic Flight**, *NAS Feature*, March 16, 2021—Researchers in the NASA Advanced Supercomputing Division are using in-house-developed tools to help the agency develop its innovative low-boom supersonic X-plane.
https://www.nas.nasa.gov/publications/articles/feature_X-59.html

News and Events: Social Media

- **Coverage of NAS Stories**
  - NAS X-59 Supersonic Low-Boom Feature:
    - NAS: Twitter 3 retweets, 14 favorites
    - NASA Supercomputing: Facebook 368 users reached, 27 engagements, 13 likes, and 4 shares.

HECC Utilization

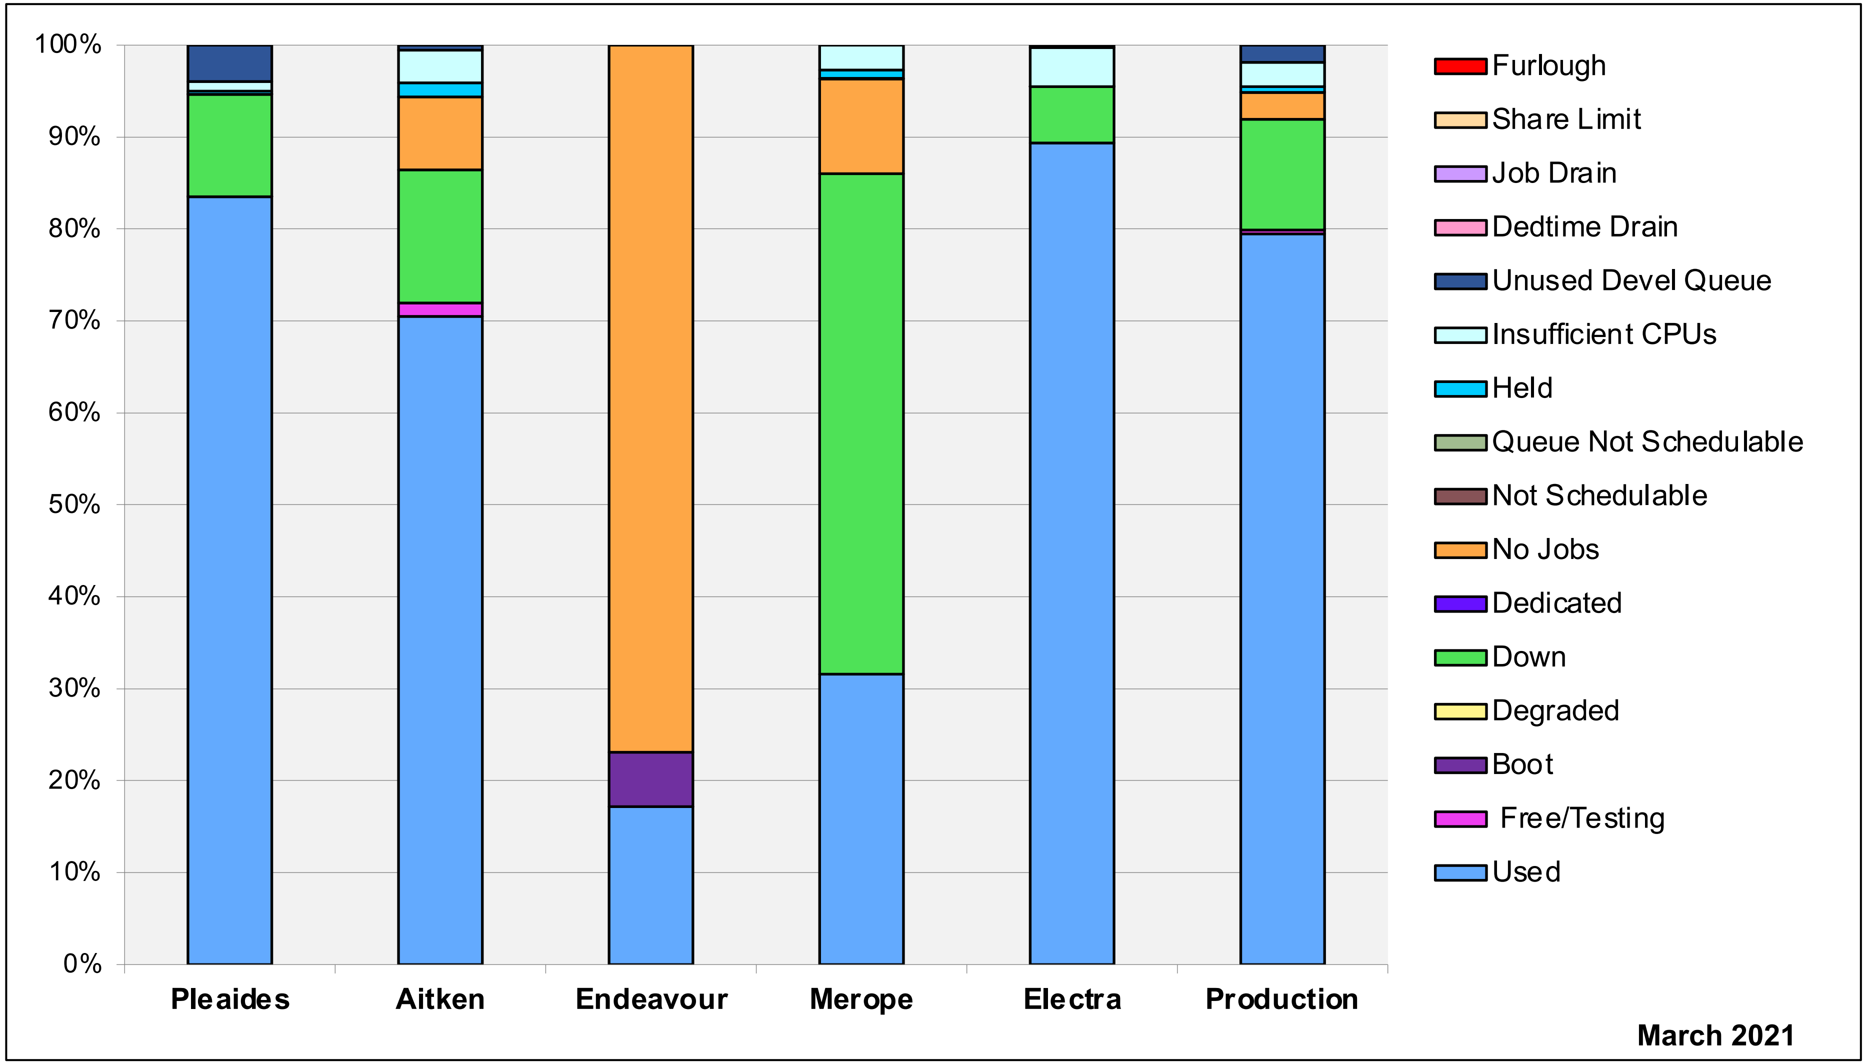

HECC Utilization Normalized to 30-Day Month

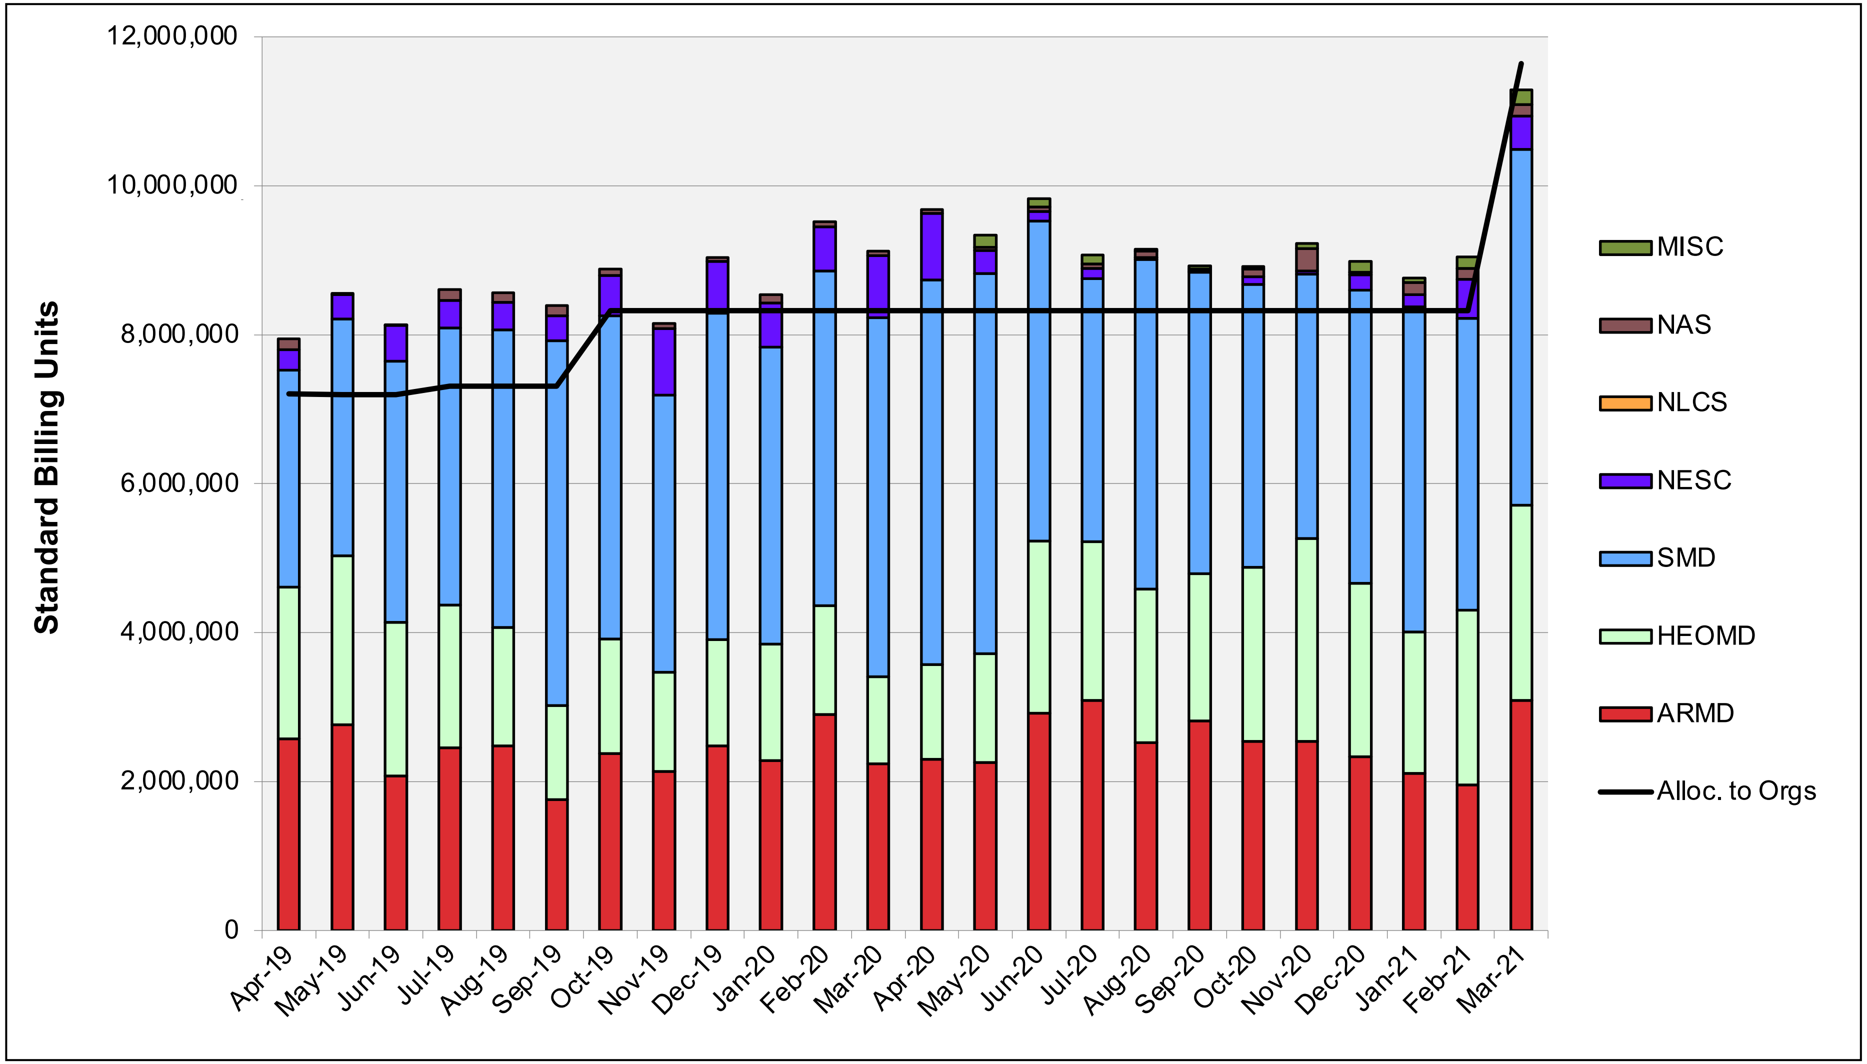

HECC Utilization Normalized to 30-Day Month

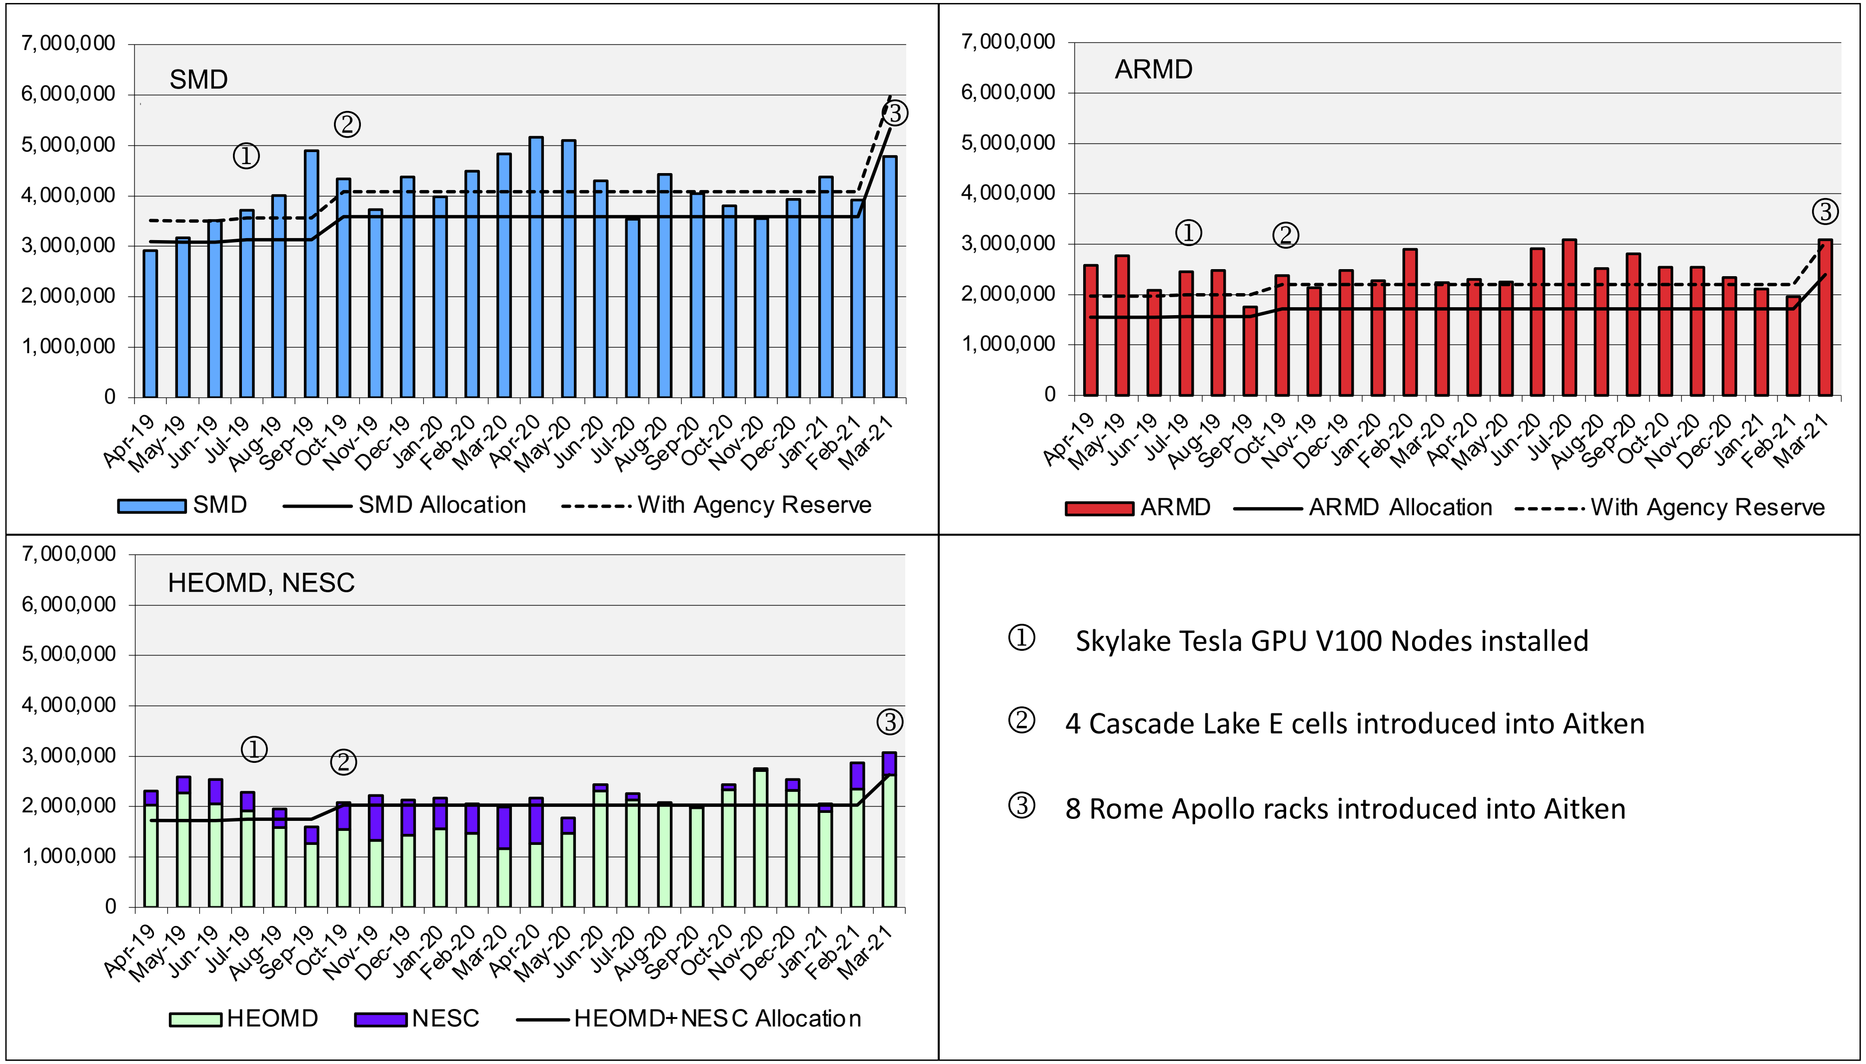

① Skylake Tesla GPU V100 Nodes installed

② 4 Cascade Lake E cells introduced into Aitken

③ 8 Rome Apollo racks introduced into Aitken

Tape Archive Status

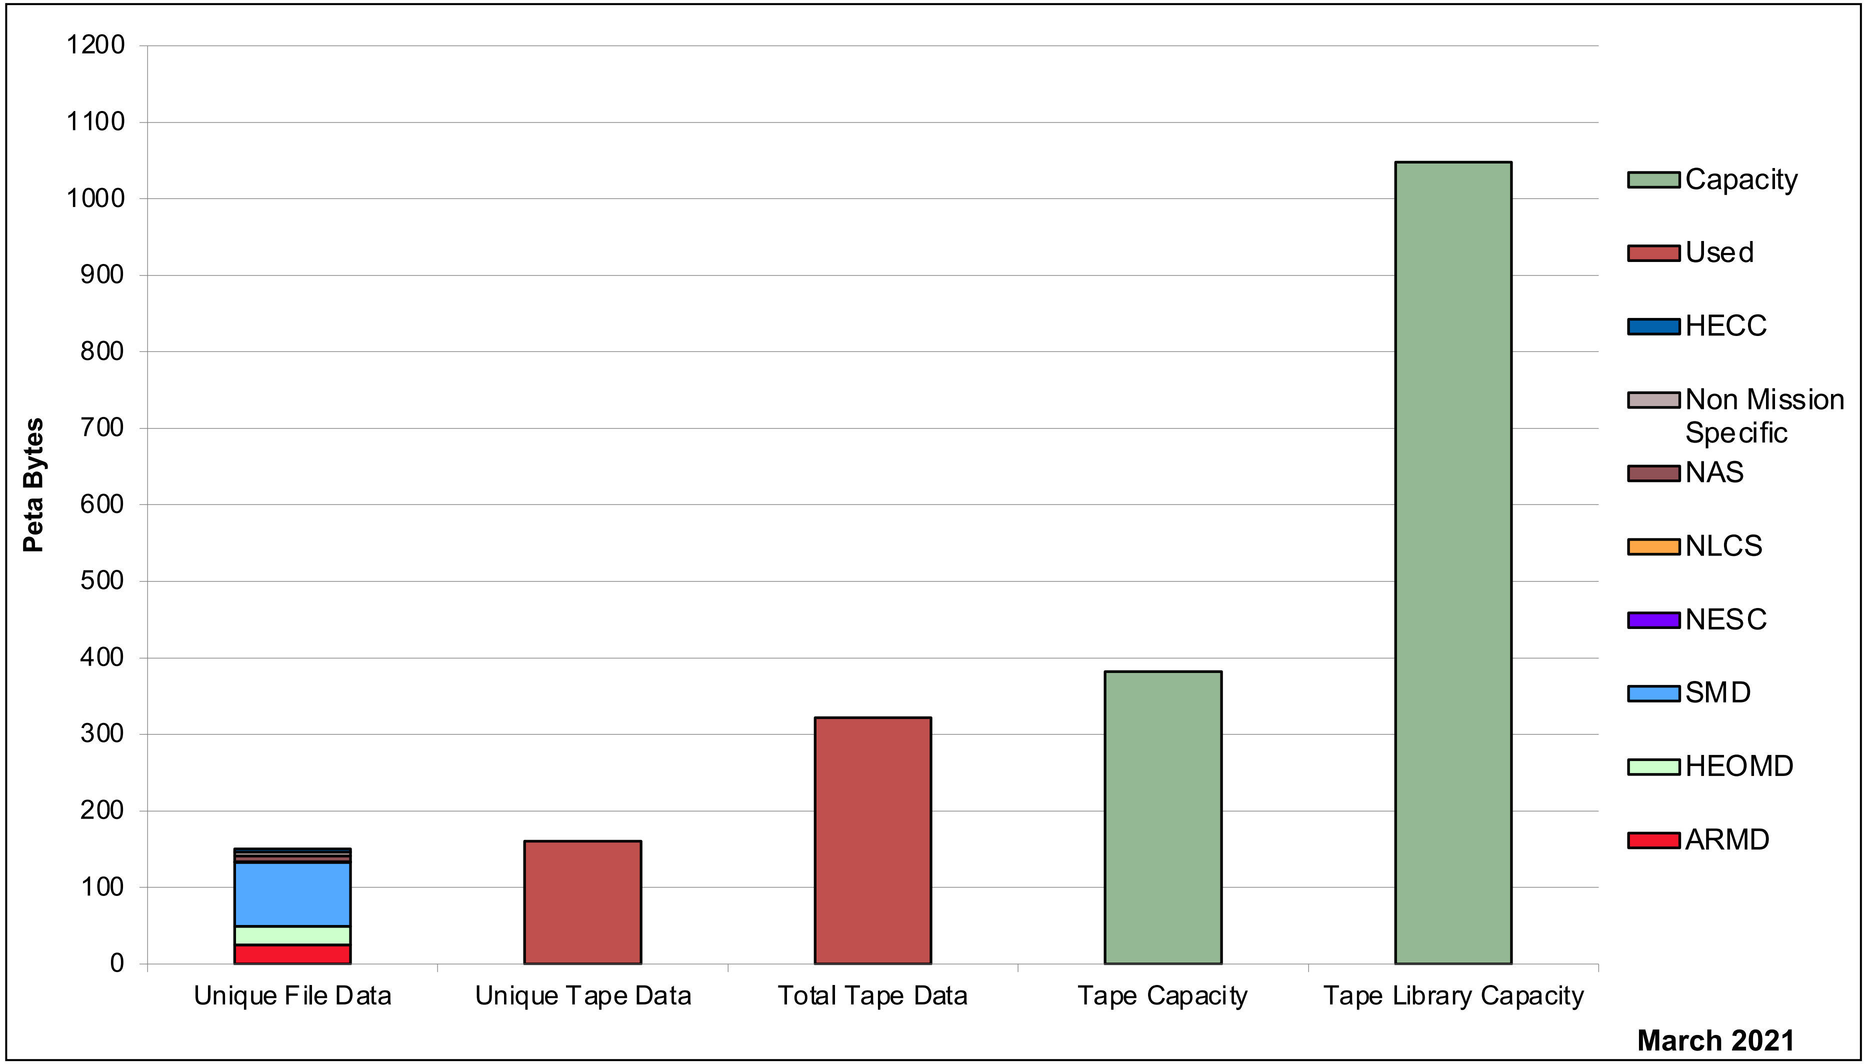

Tape Archive Status

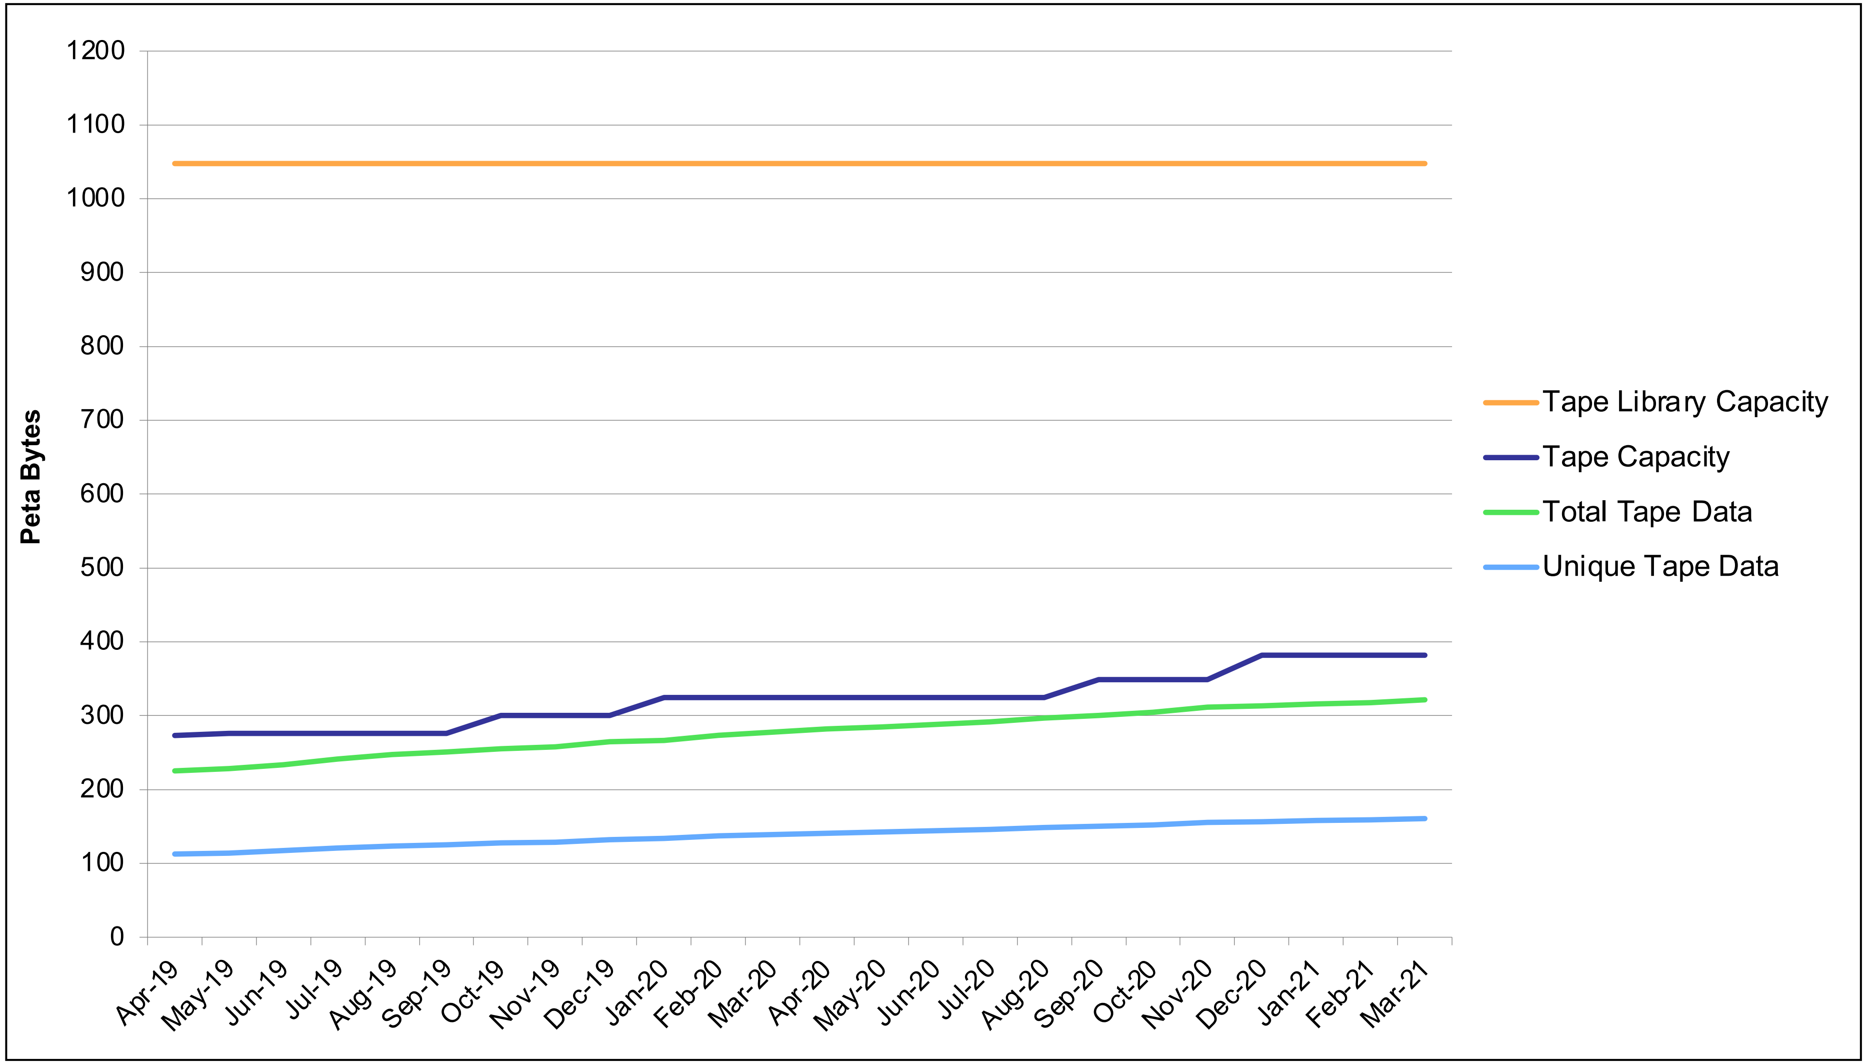

Pleiades: SBUs Reported, Normalized to 30-Day Month

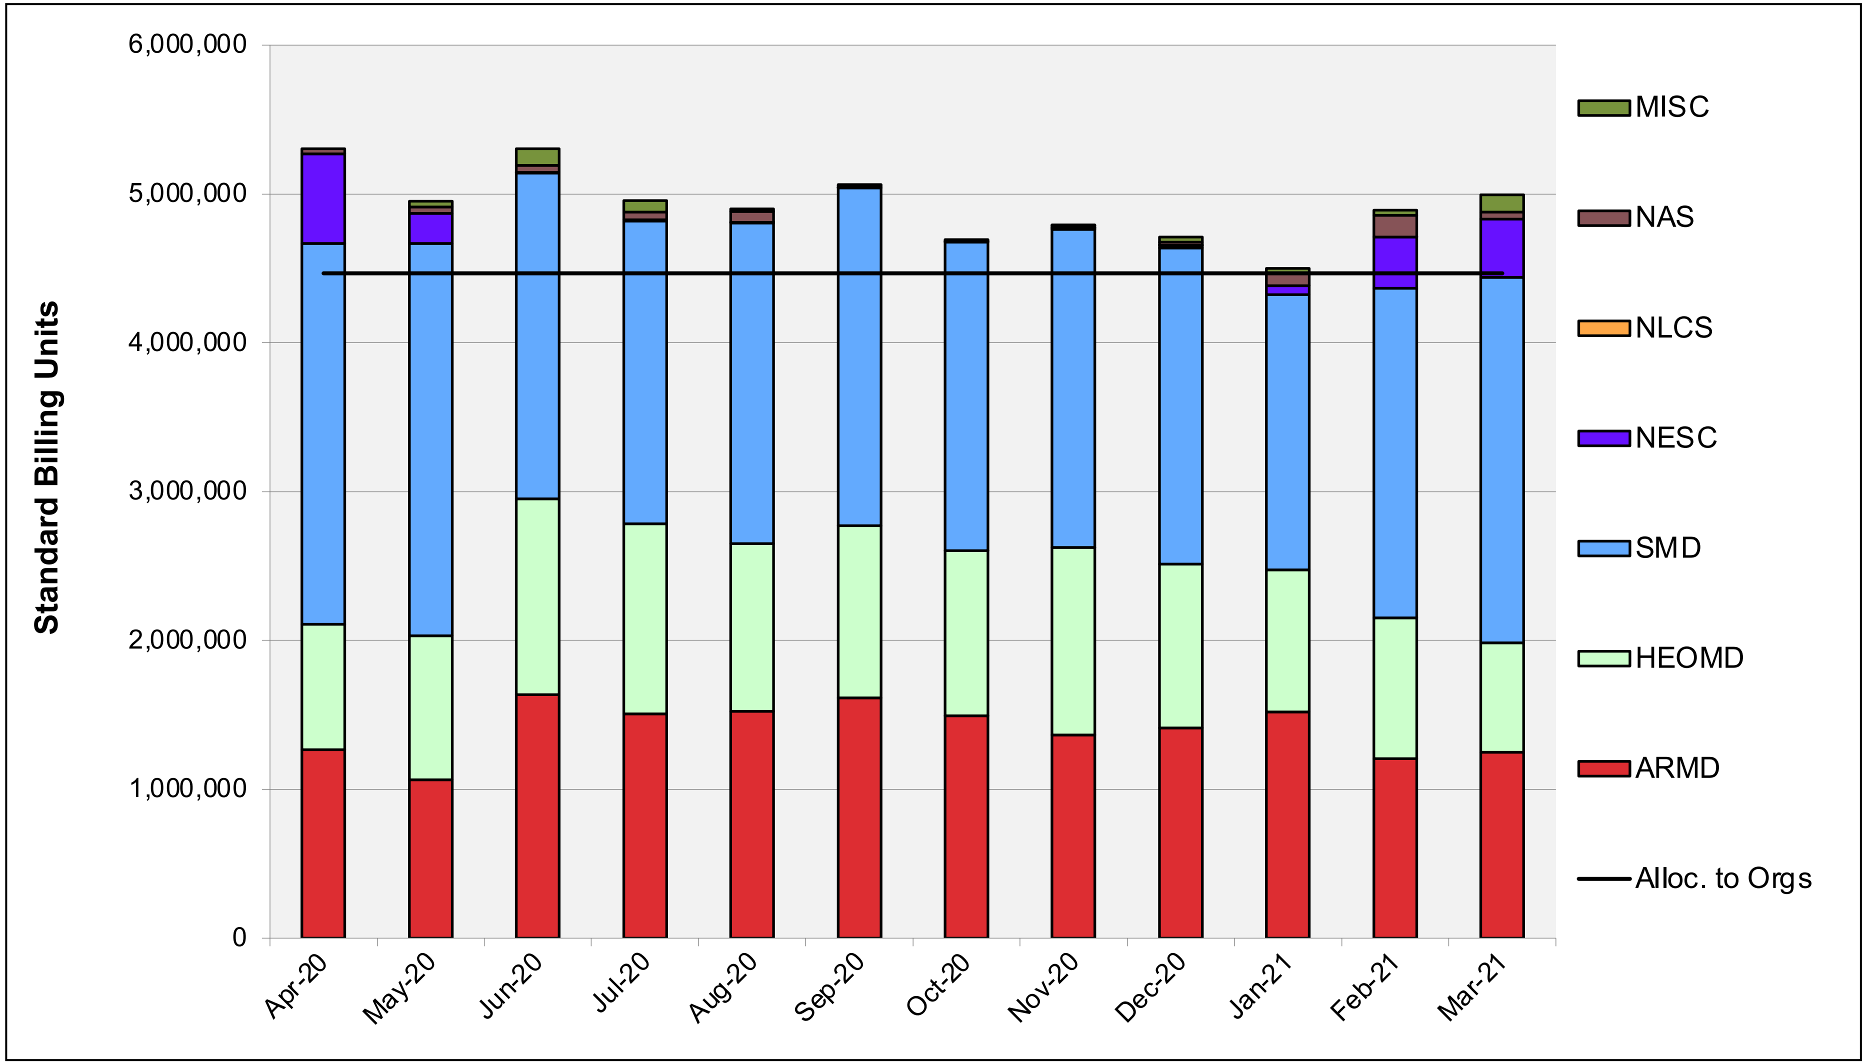

Pleiades: Devel Queue Utilization

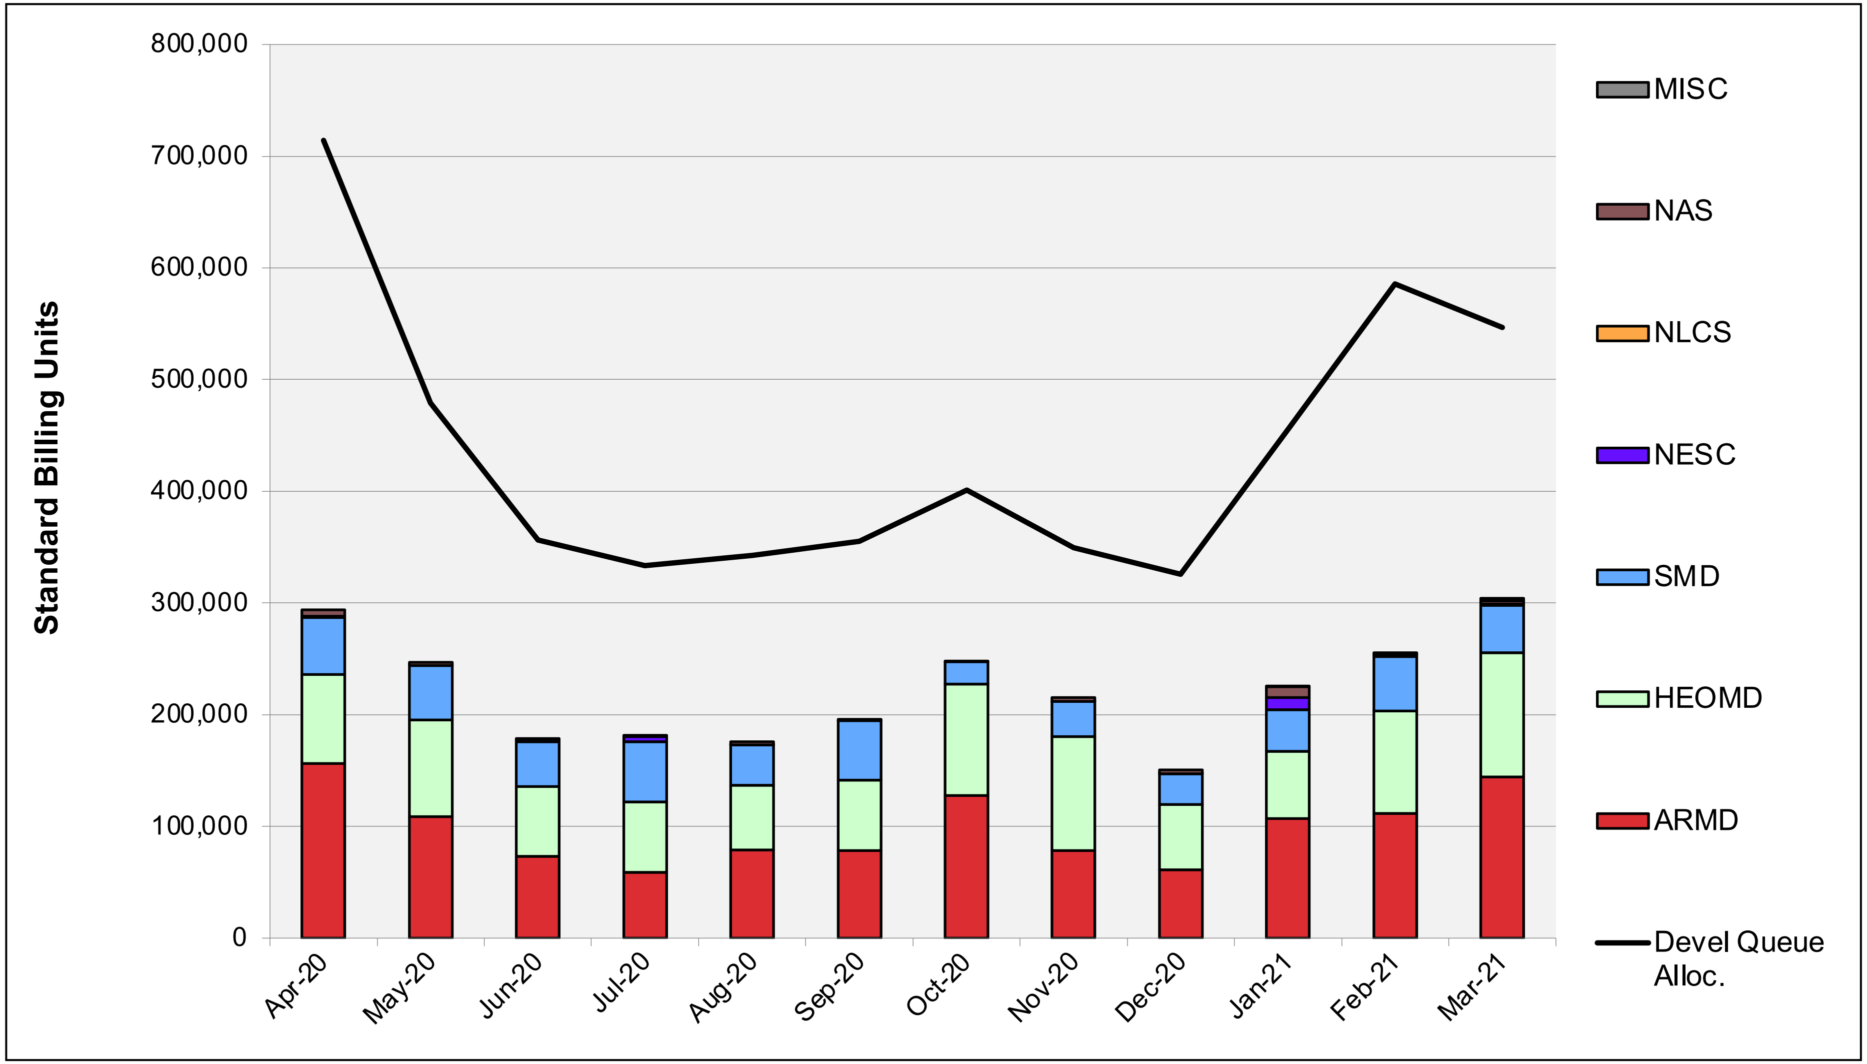

Pleiades: Monthly Utilization by Job Length

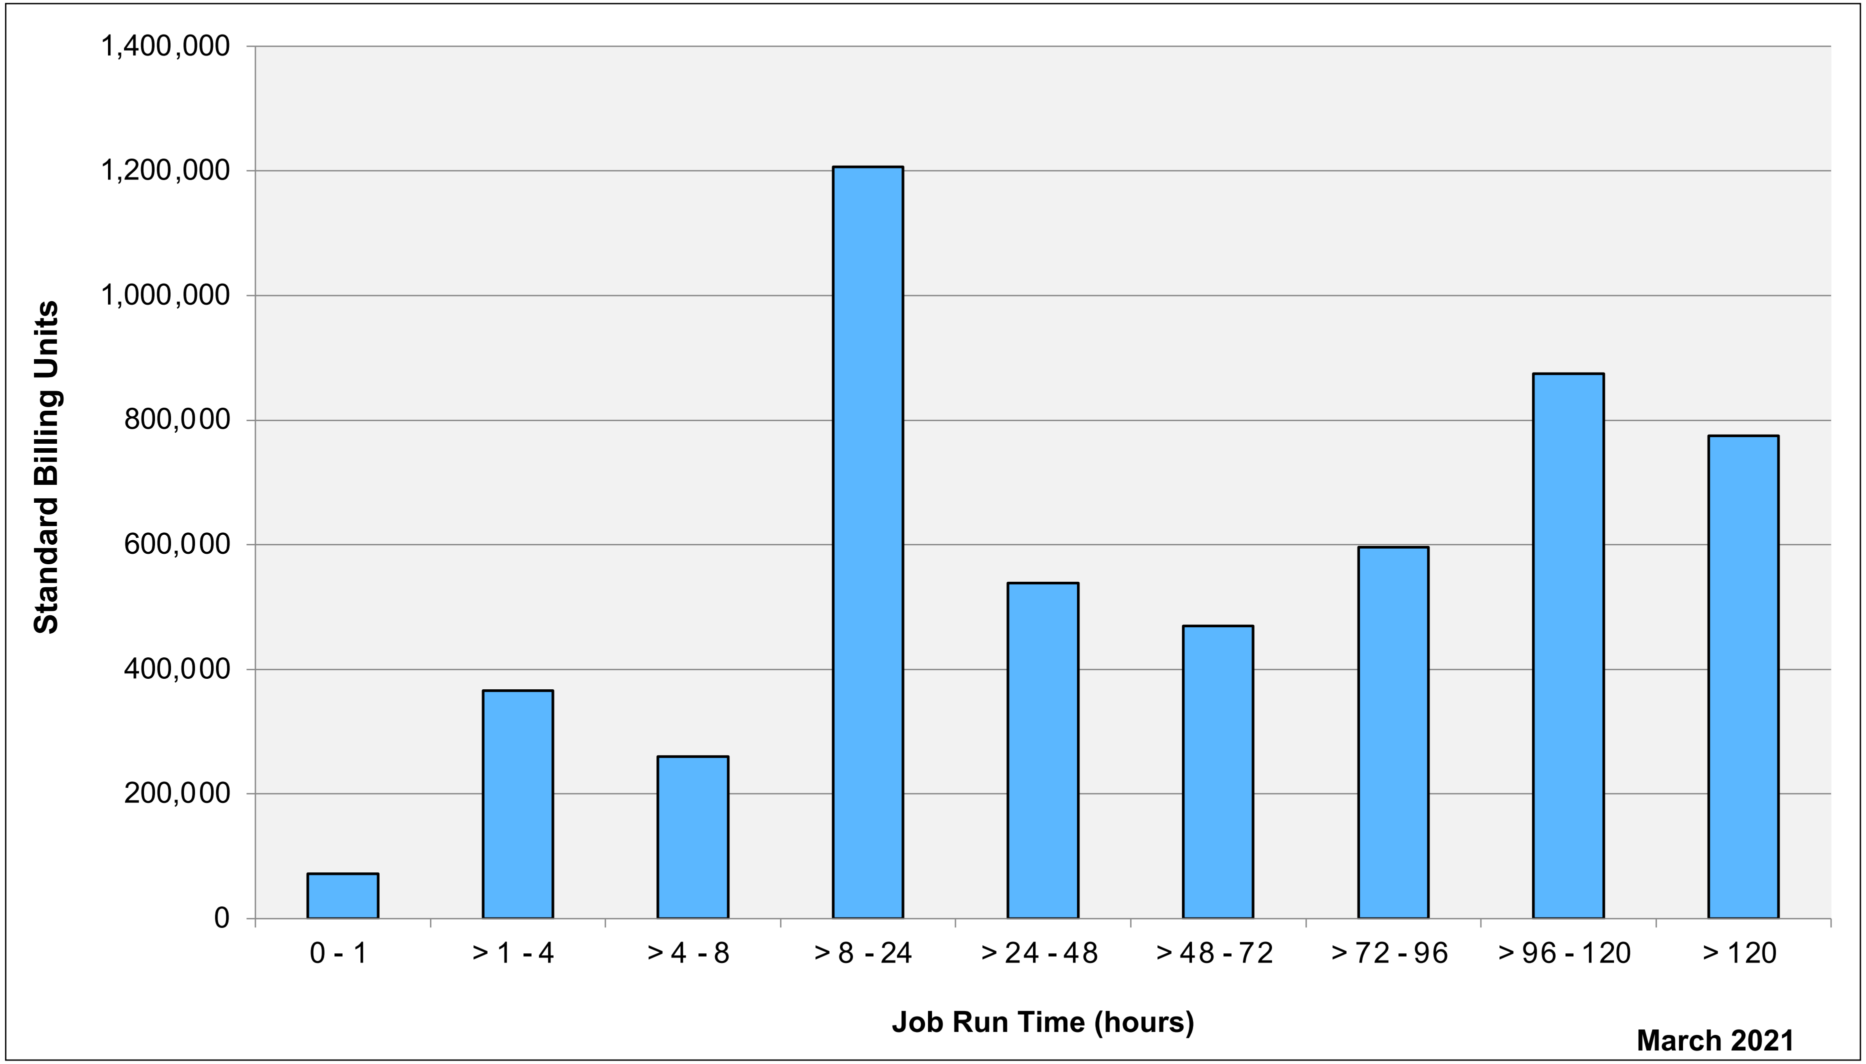

Pleiades: Monthly Utilization by Job Size

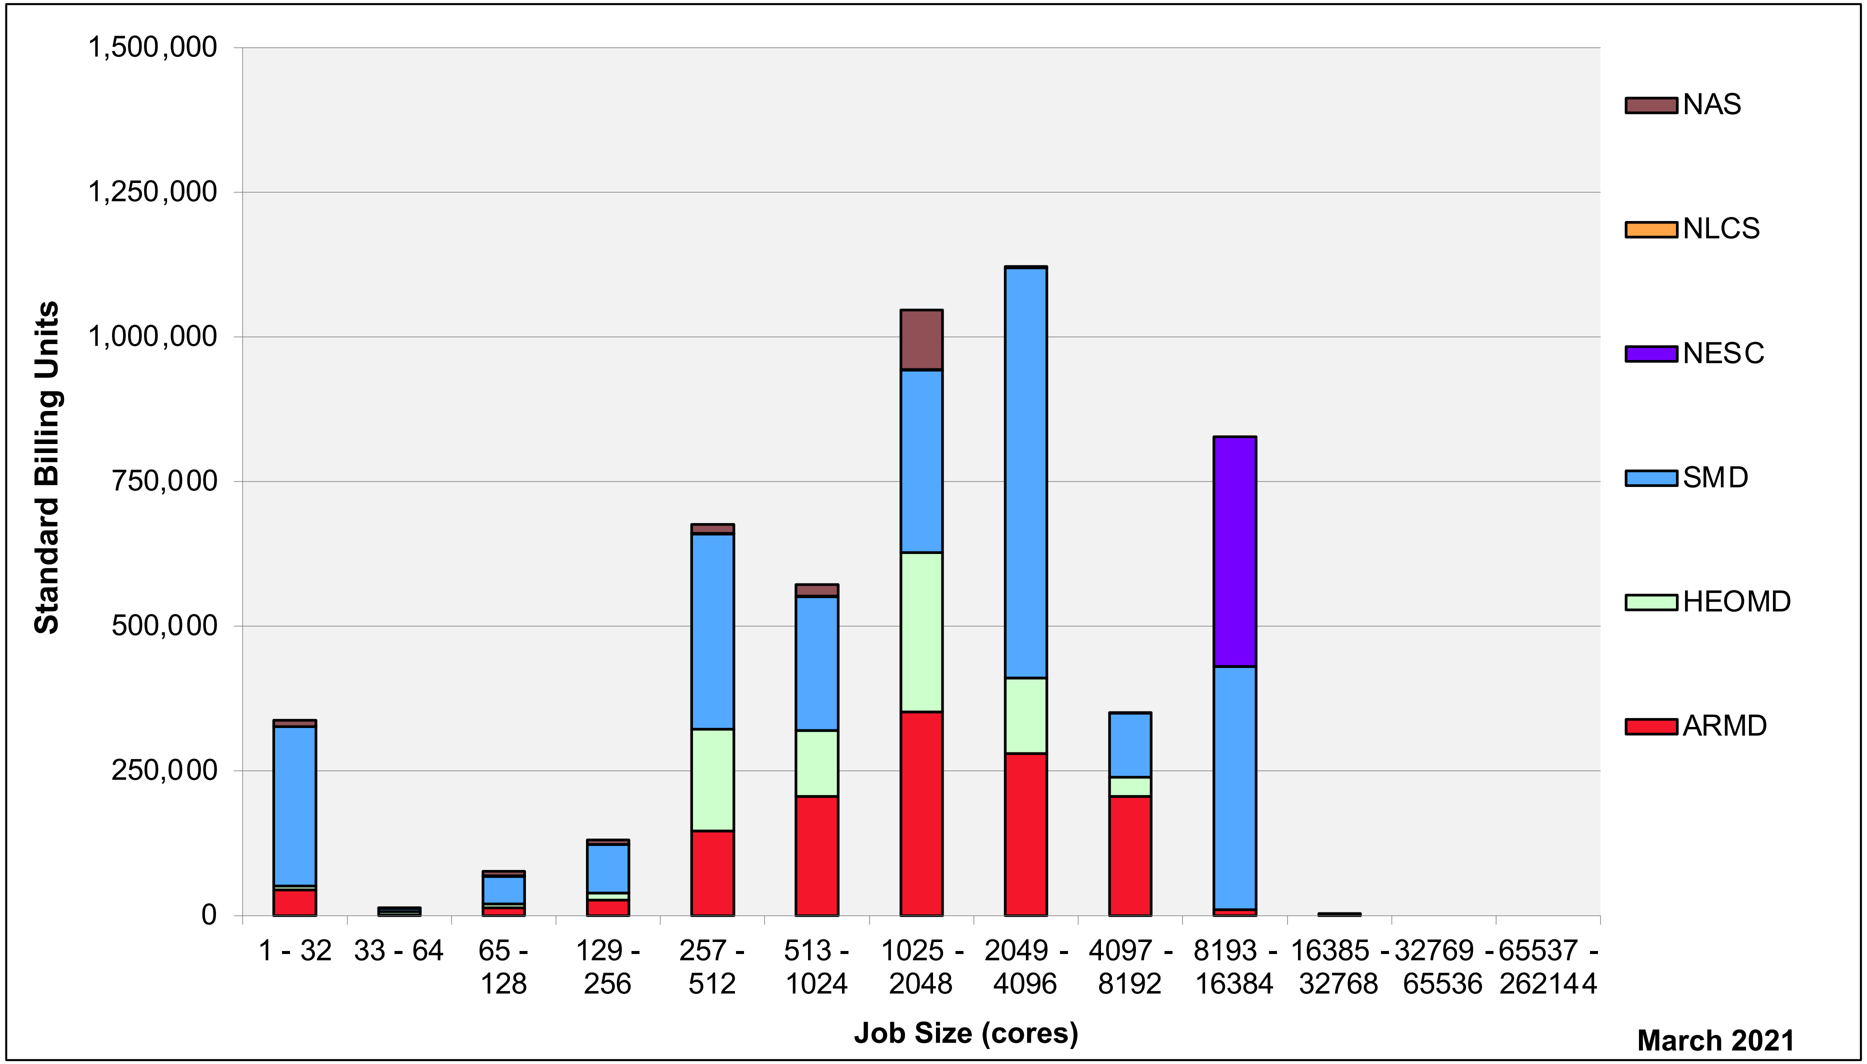

Pleiades: Monthly Utilization by Size and Length

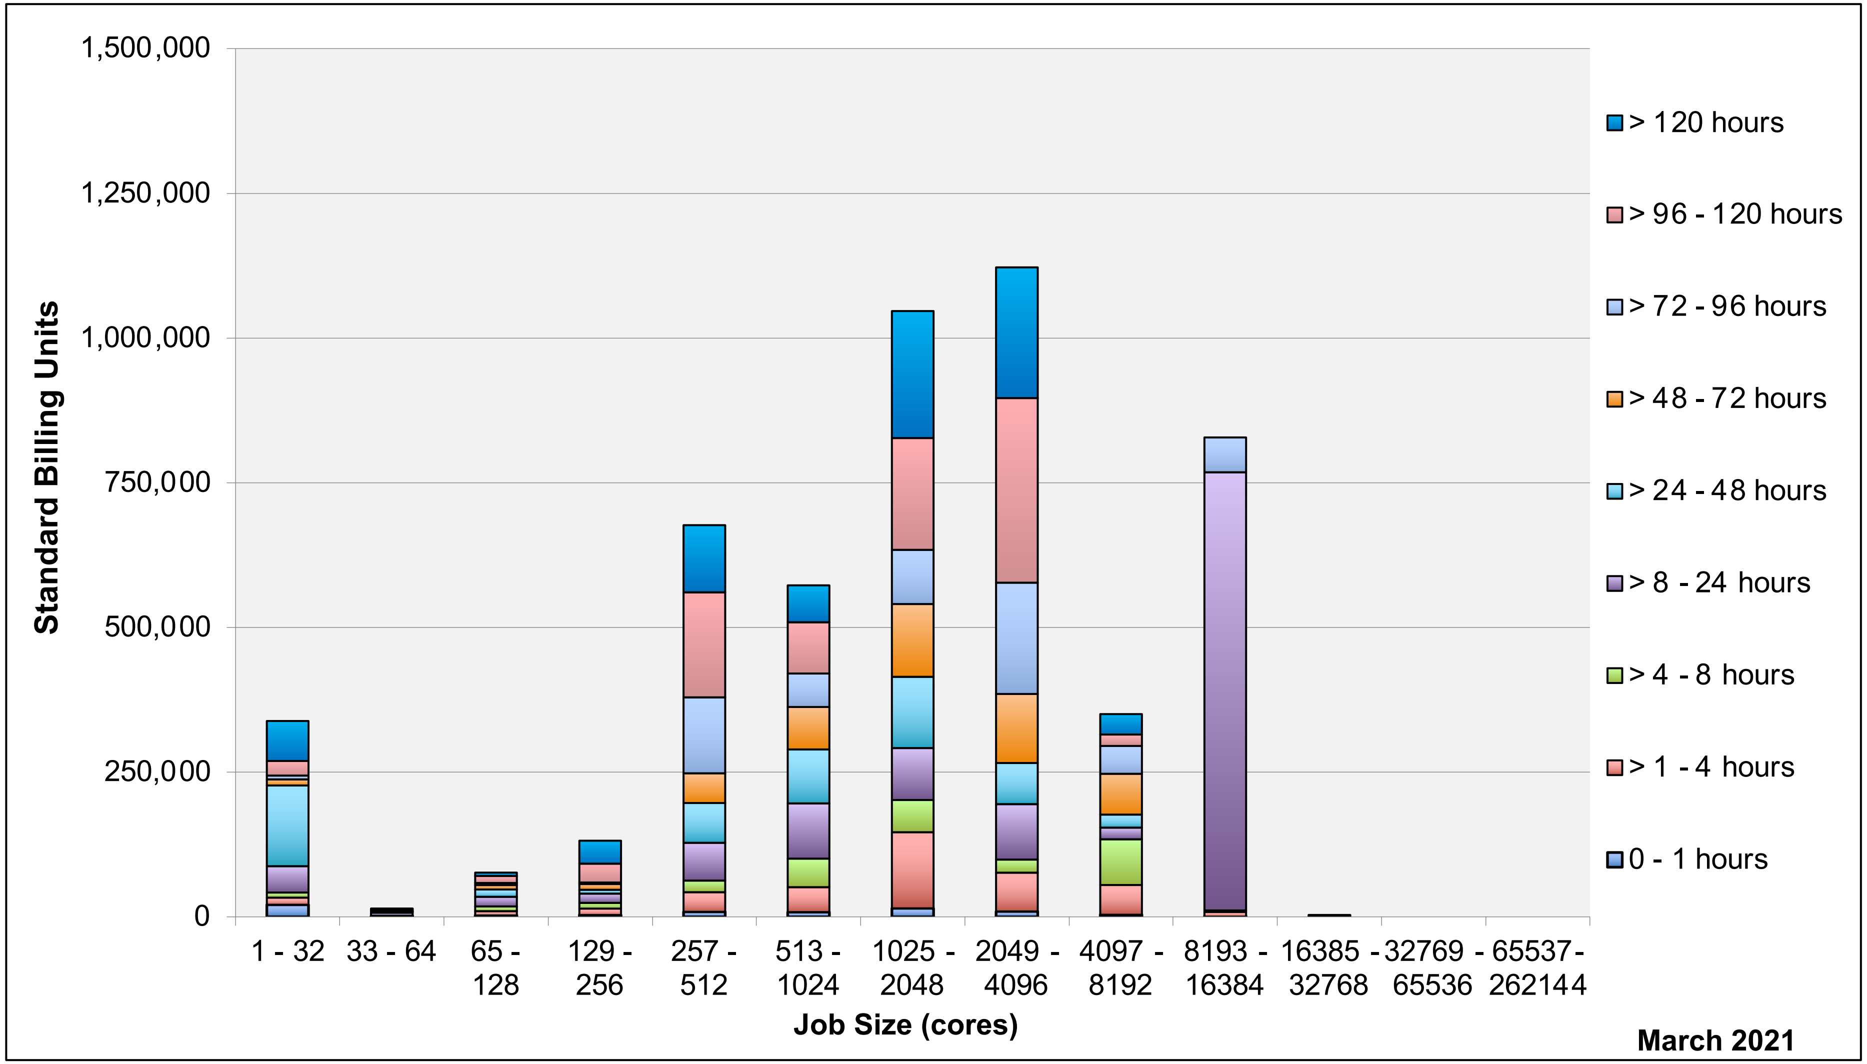

Pleiades: Average Time to Clear All Jobs

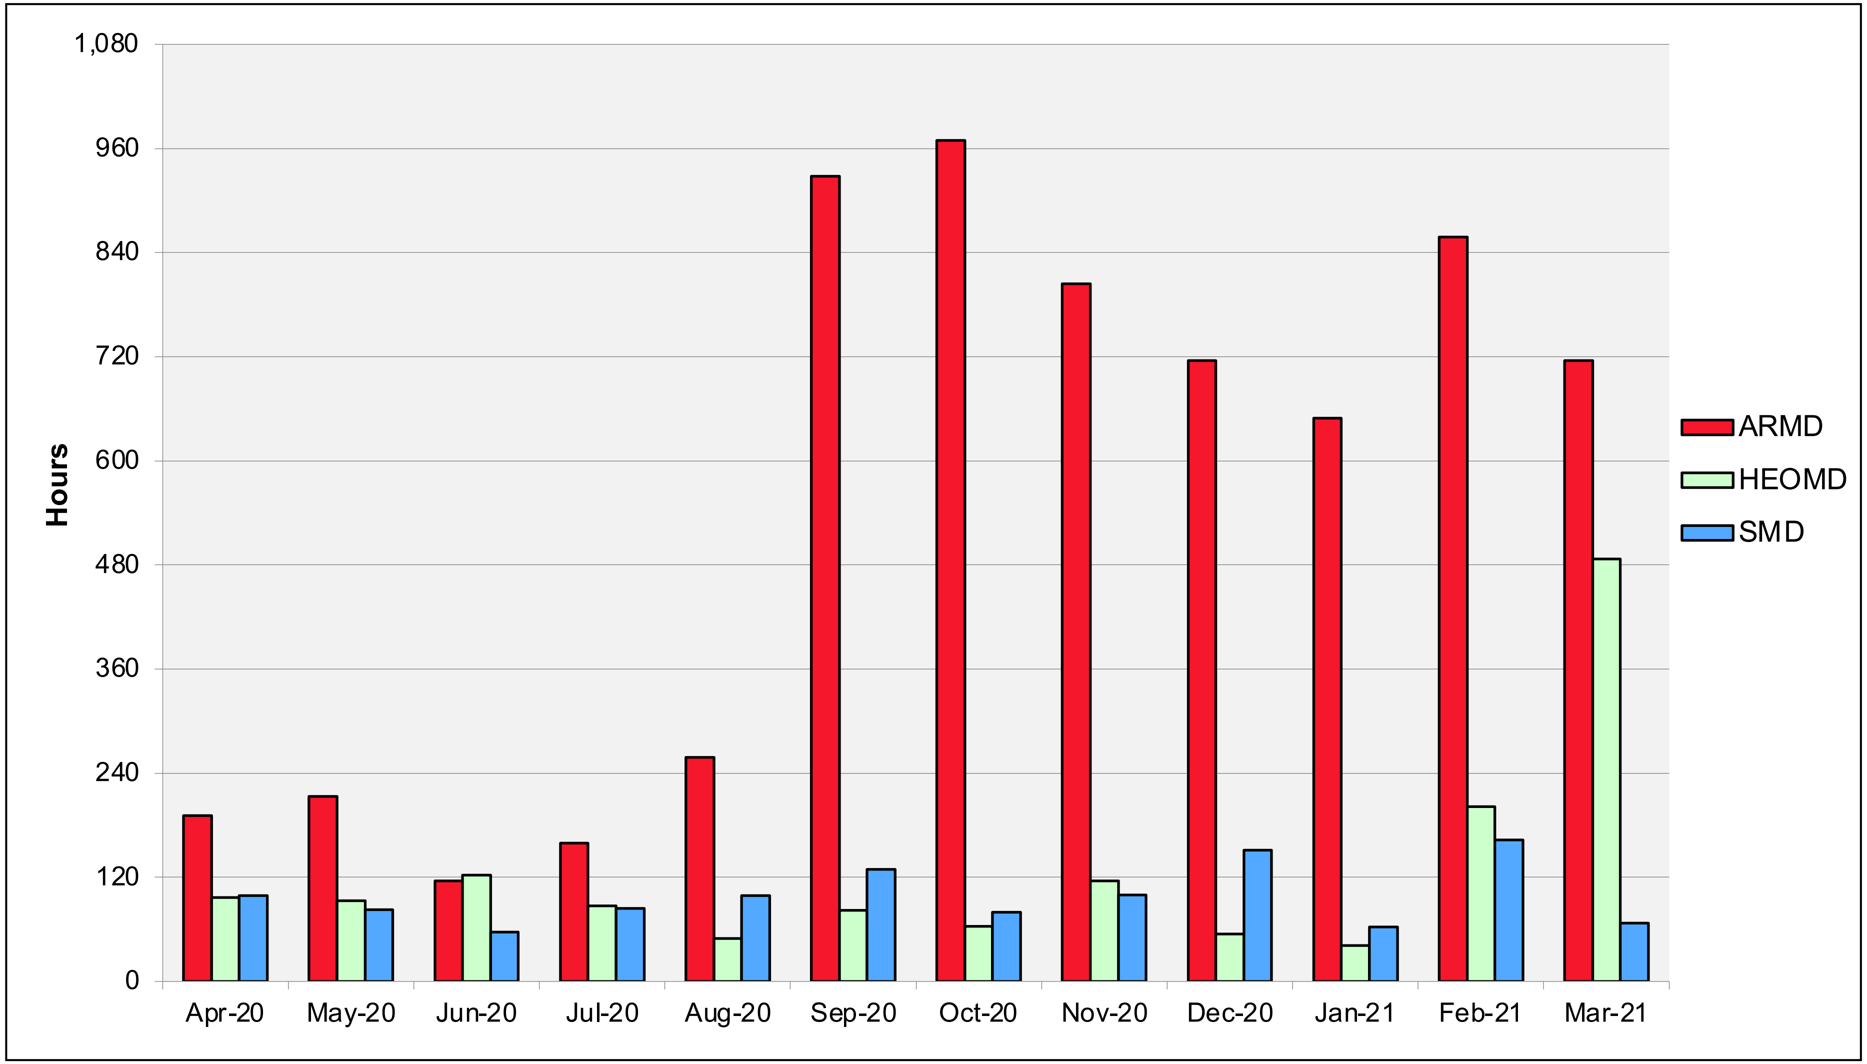

Pleiades: Average Expansion Factor

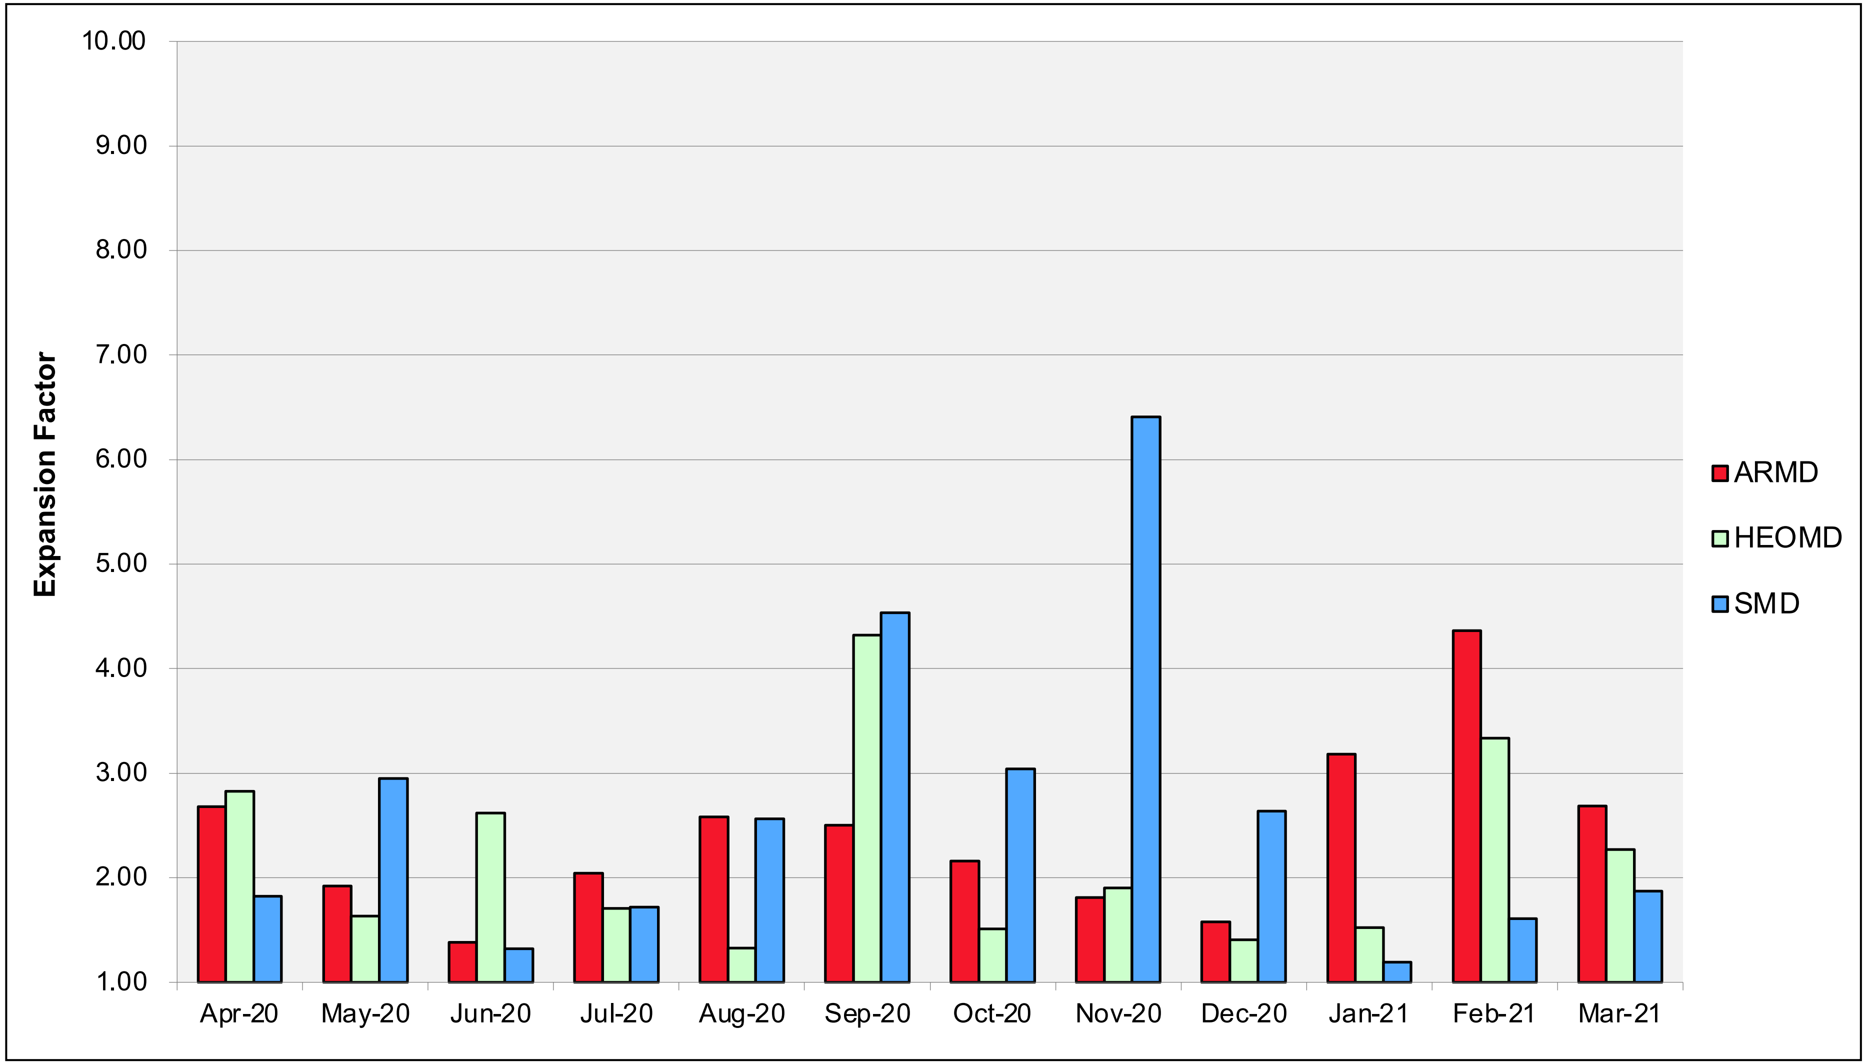

Aitken: SBUs Reported, Normalized to 30-Day Month

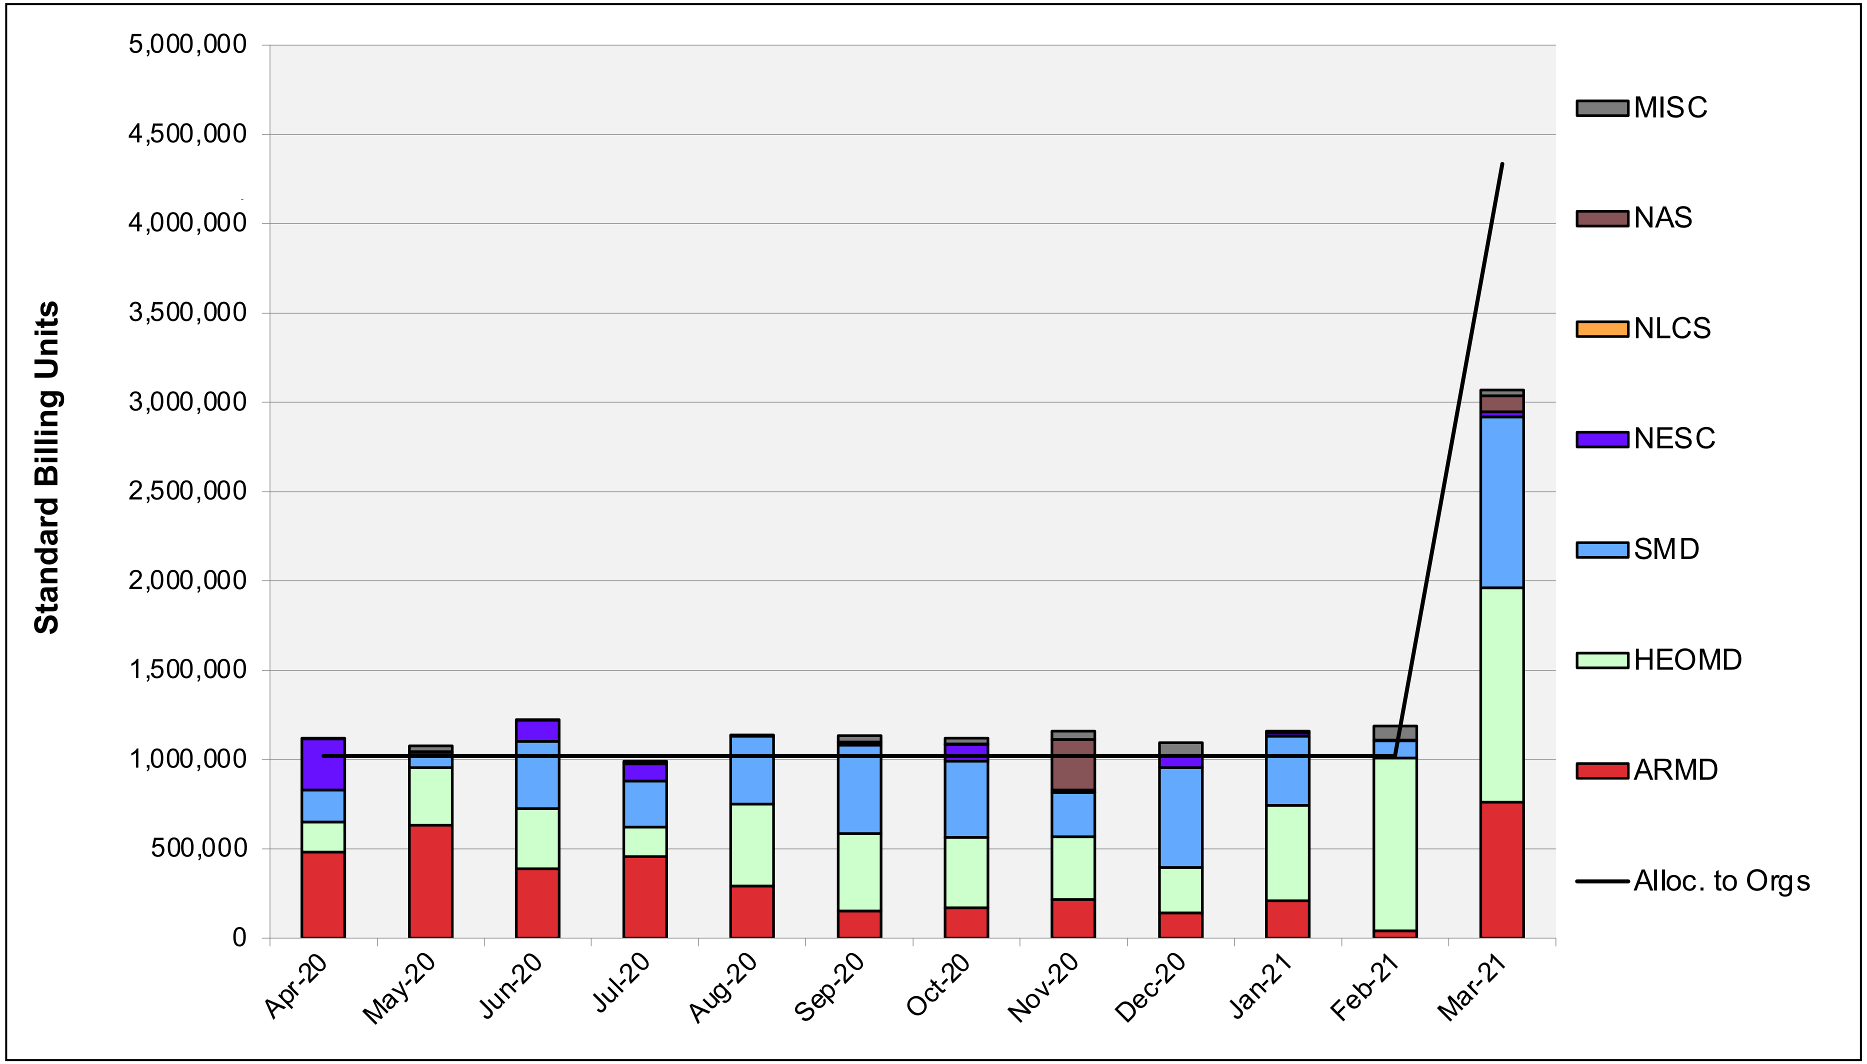

Aitken: Devel Queue Utilization

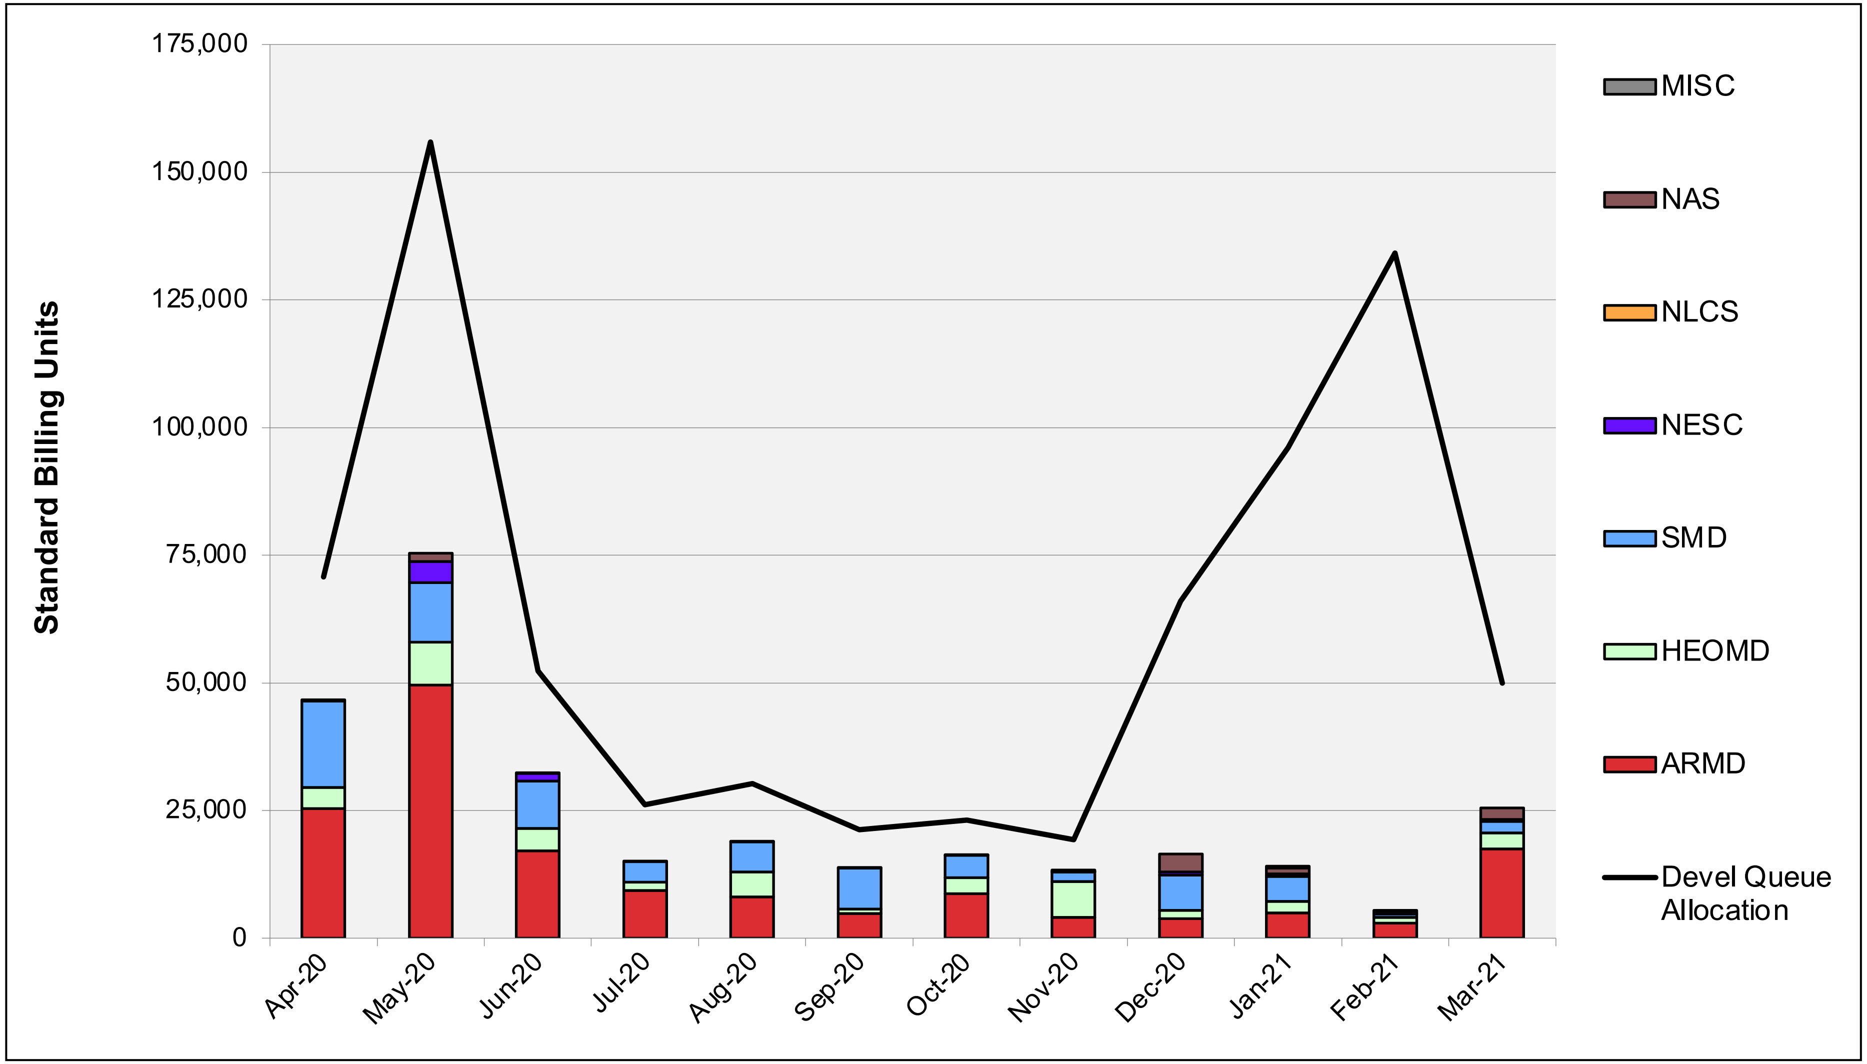

Aitken: Monthly Utilization by Job Length

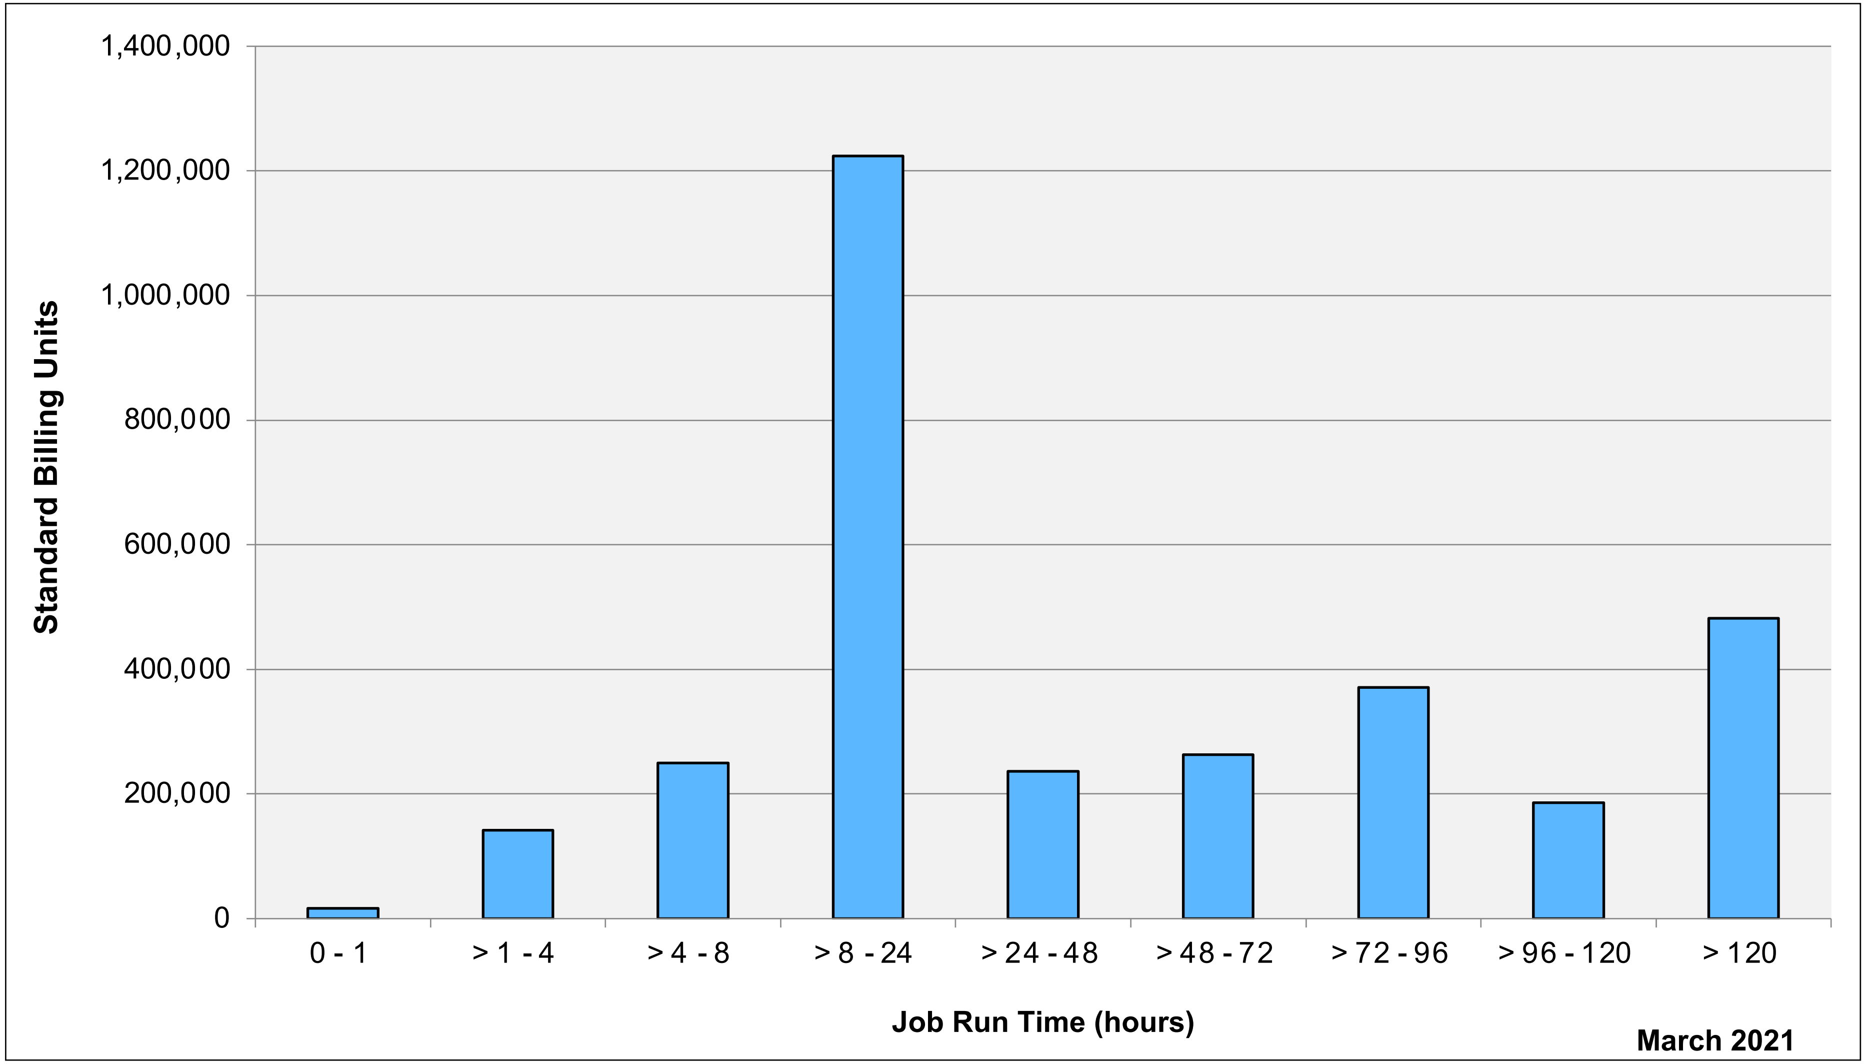

Aitken: Monthly Utilization by Job Size

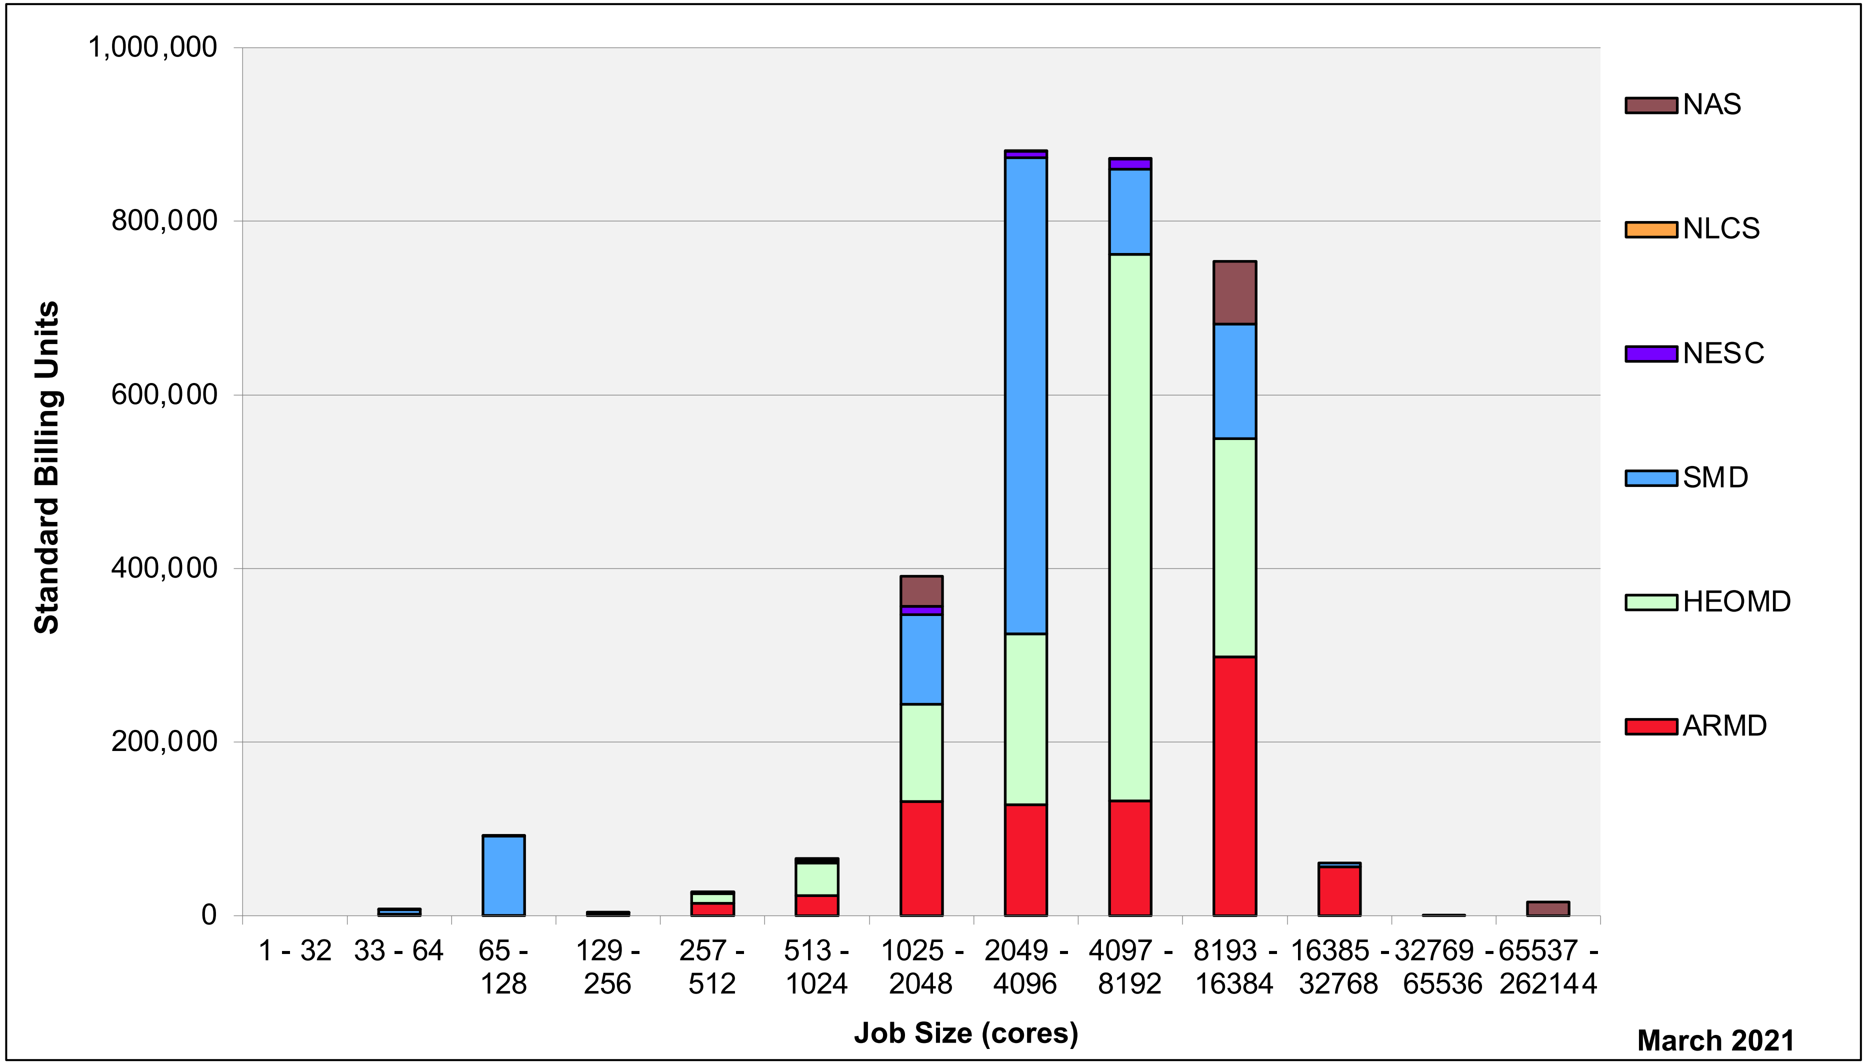

Aitken: Monthly Utilization by Size and Length

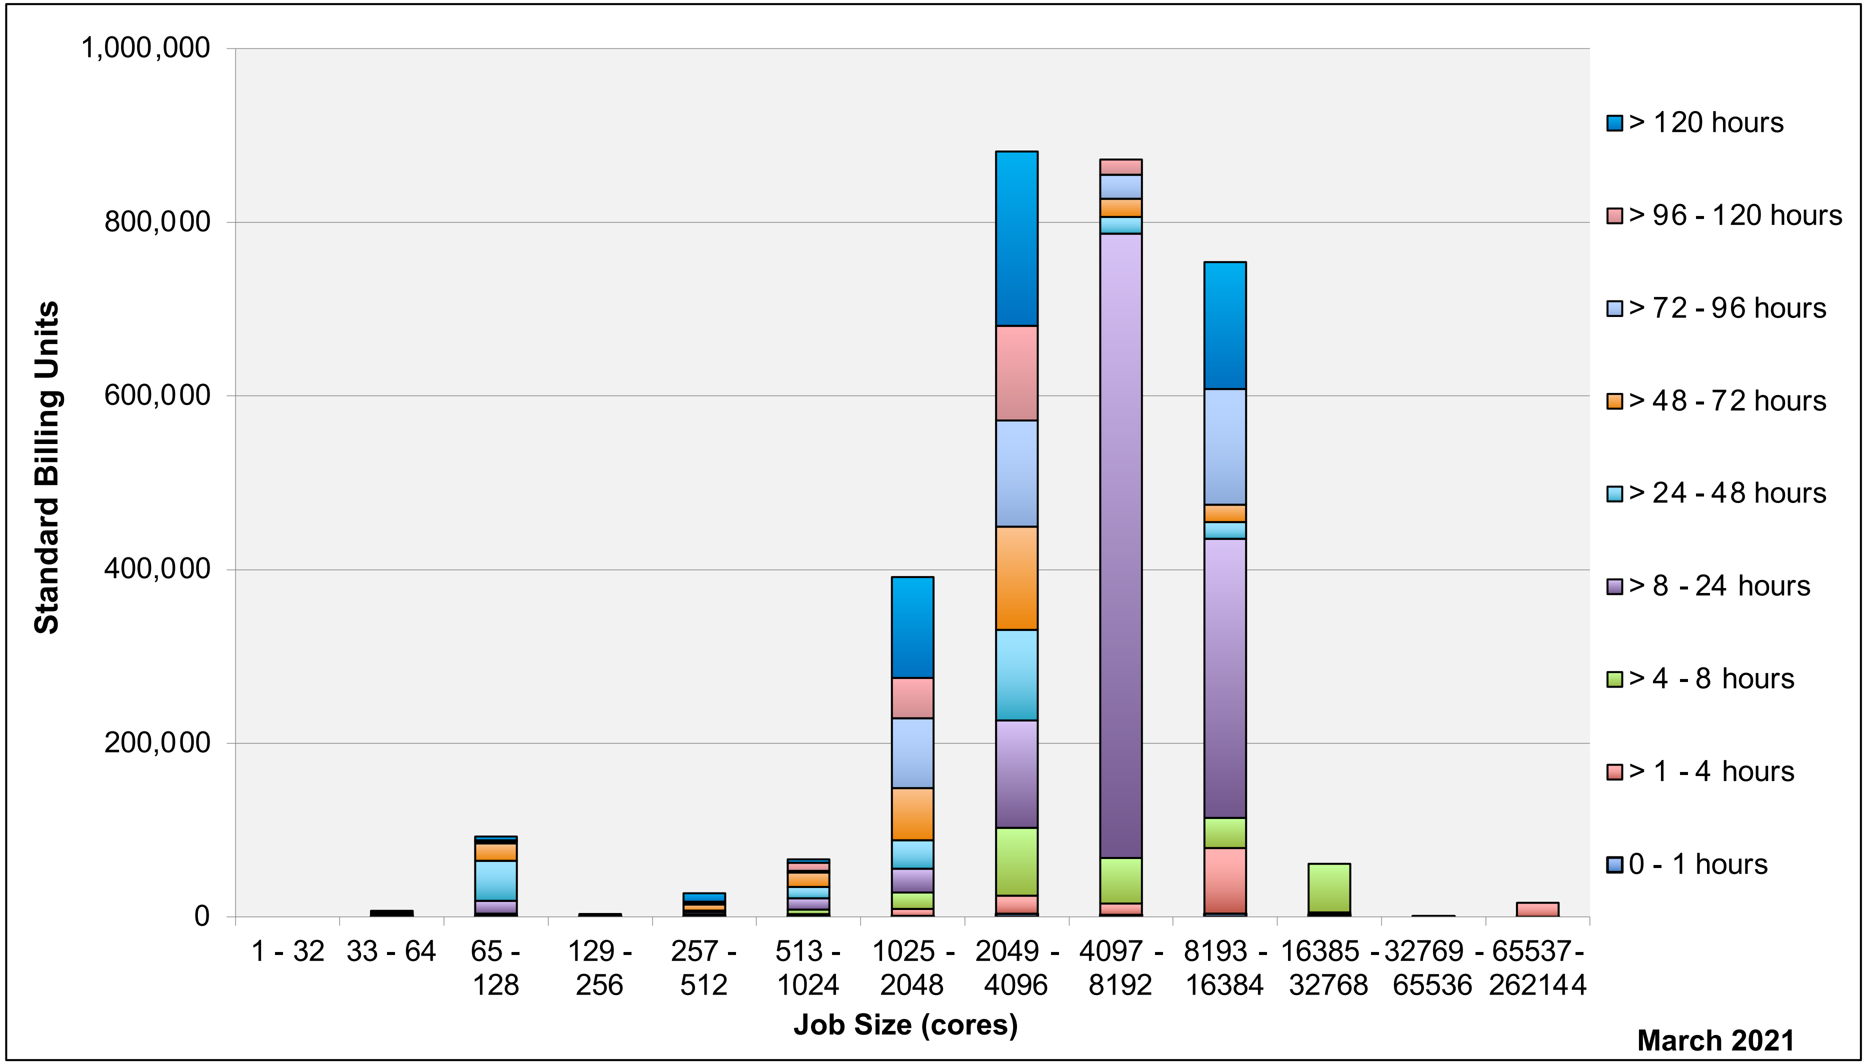

Aitken: Average Time to Clear All Jobs

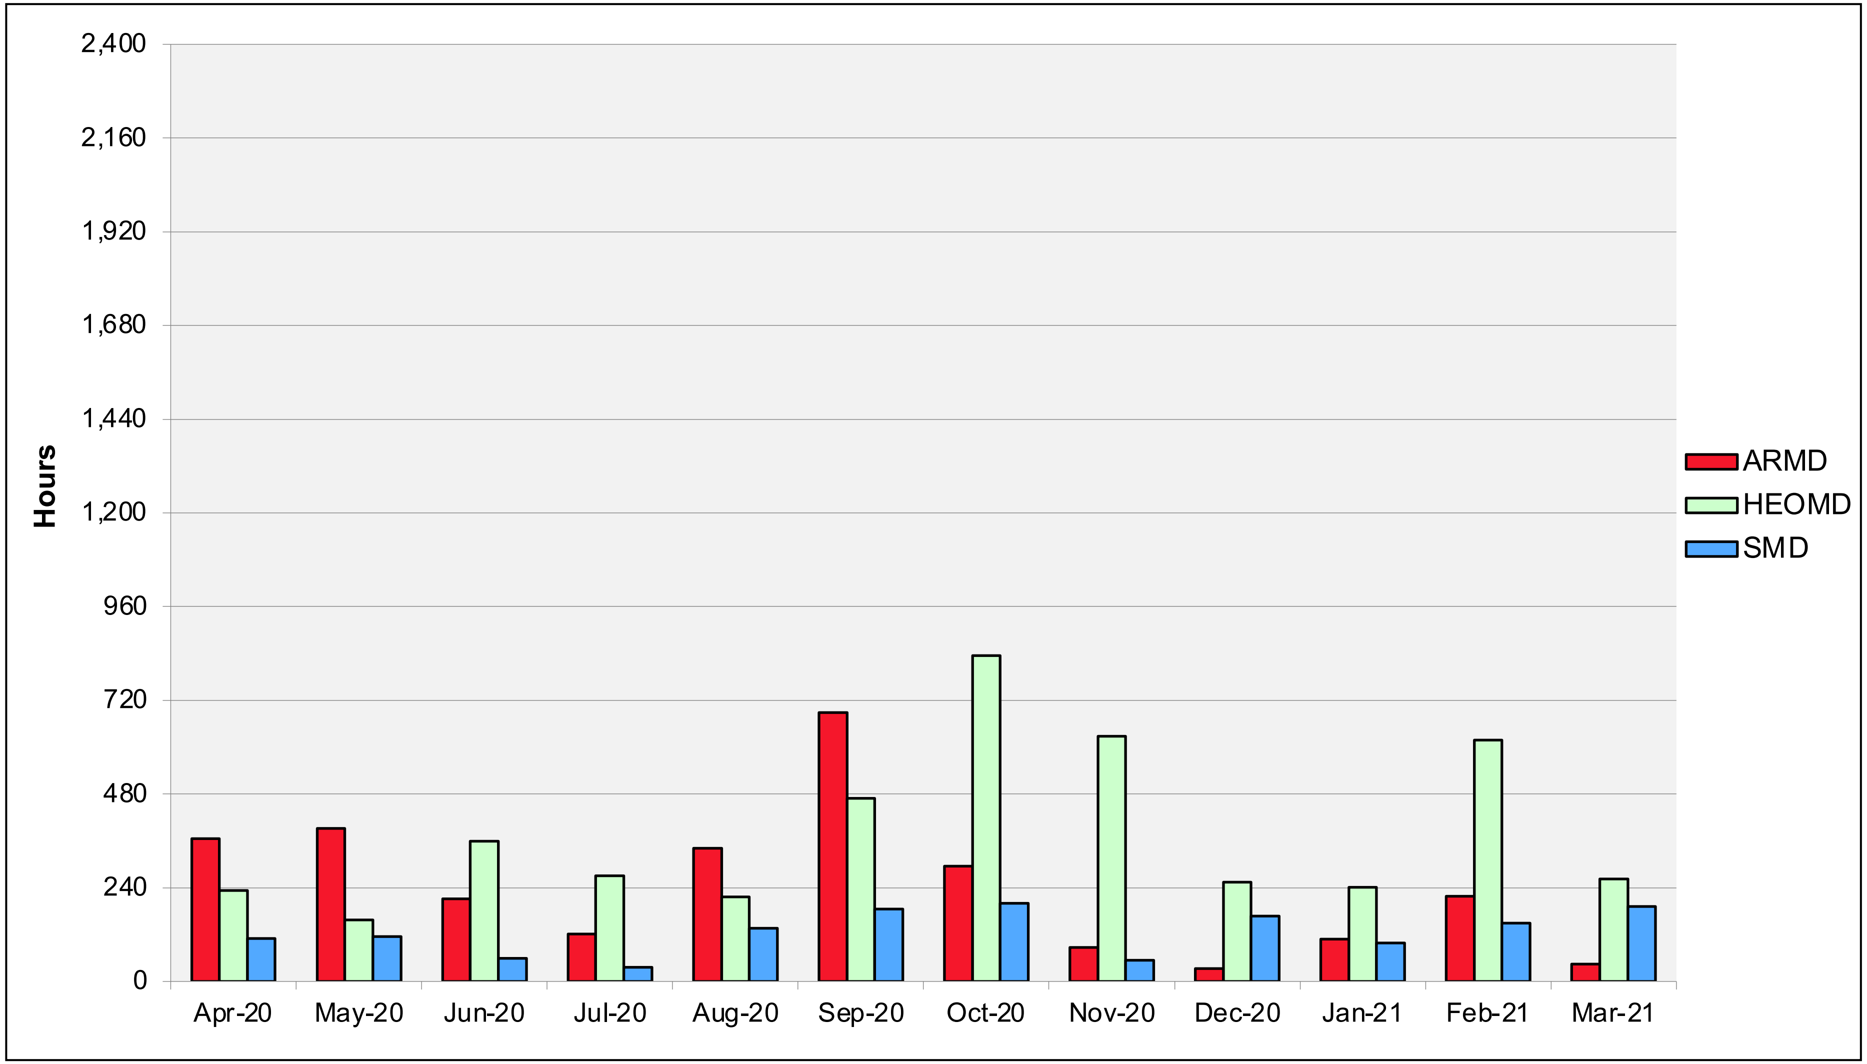

Aitken: Average Expansion Factor

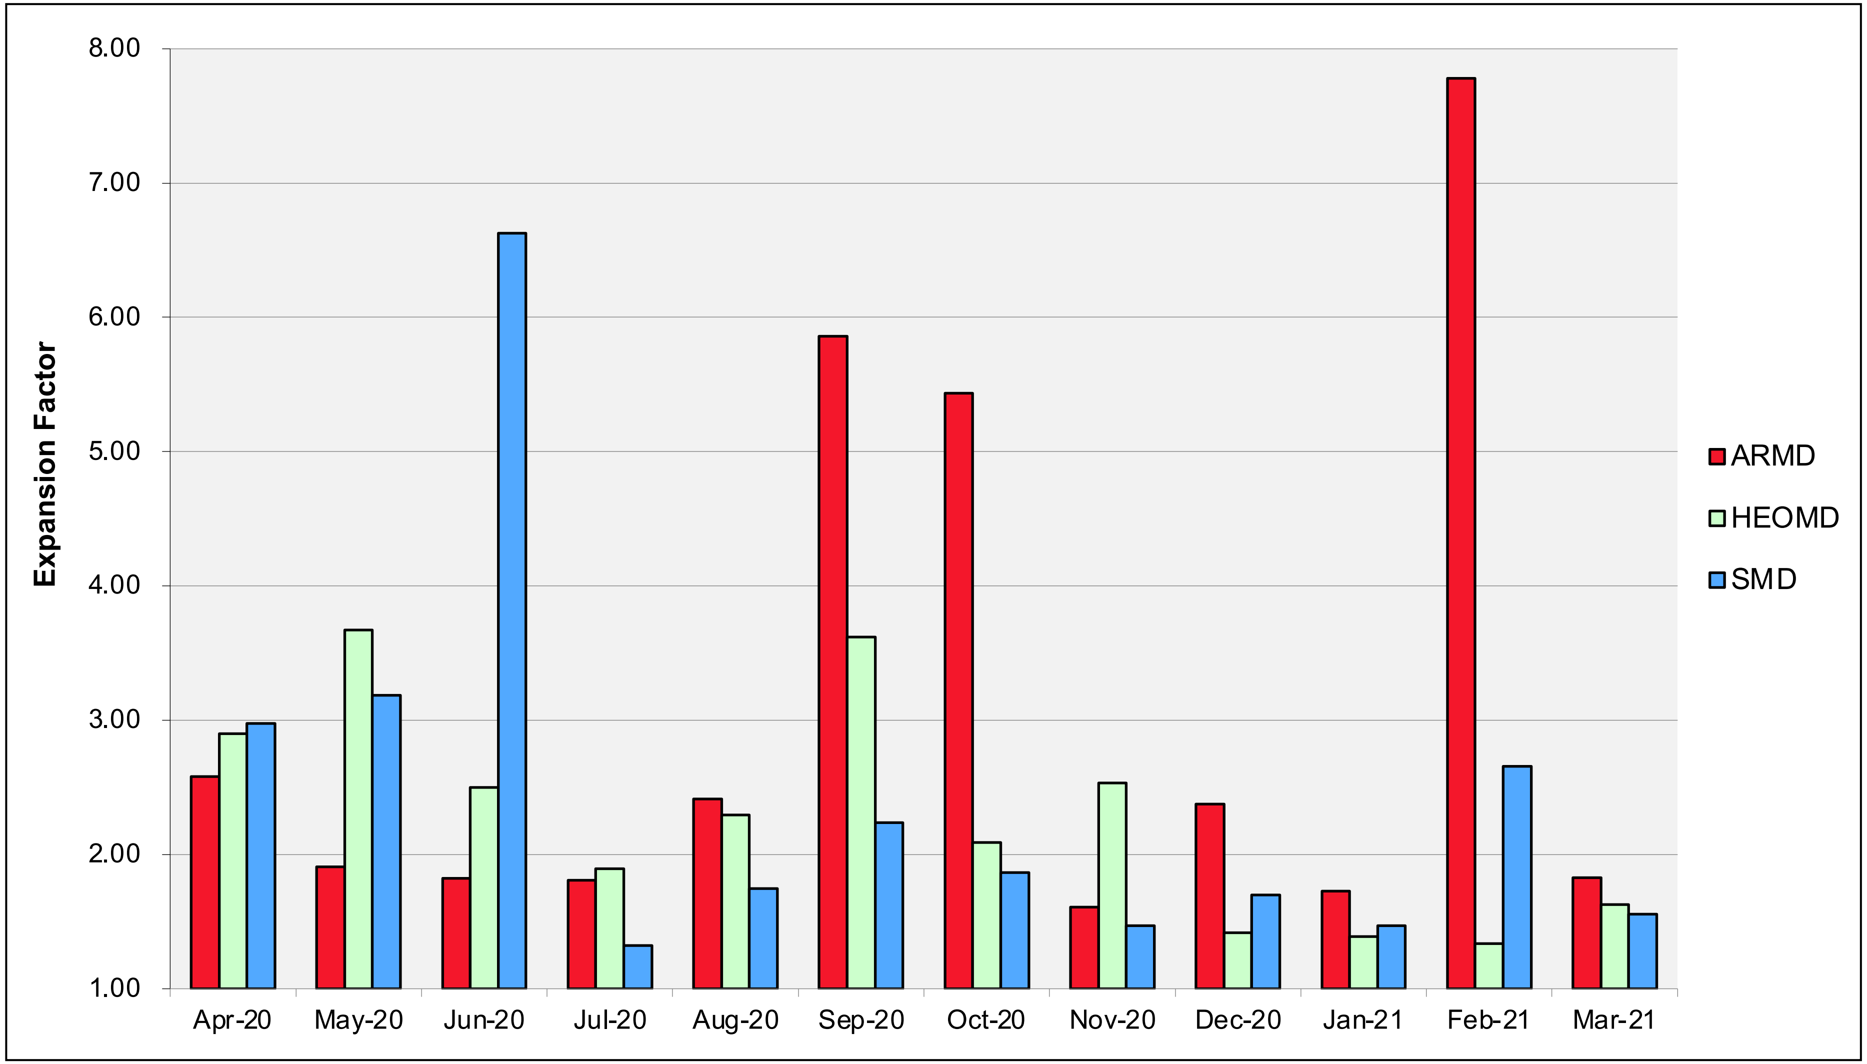

Electra: SBUs Reported, Normalized to 30-Day Month

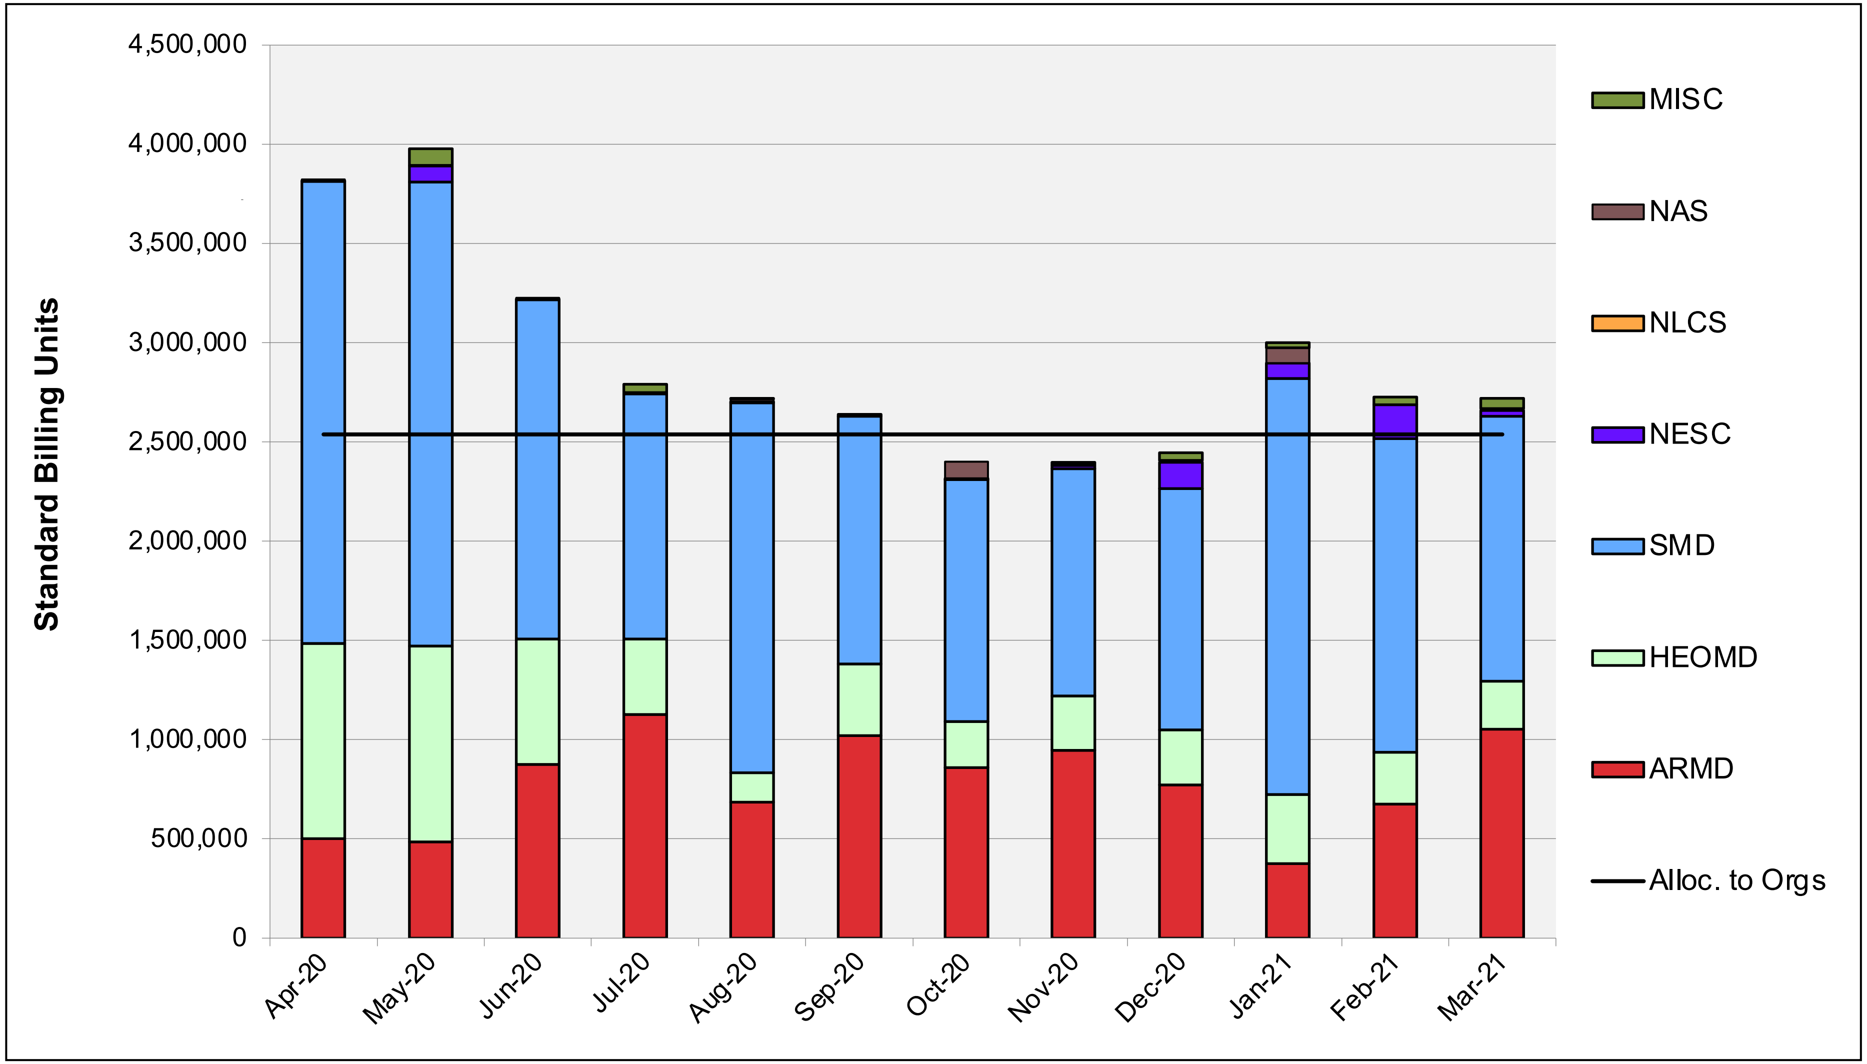

Electra: Devel Queue Utilization

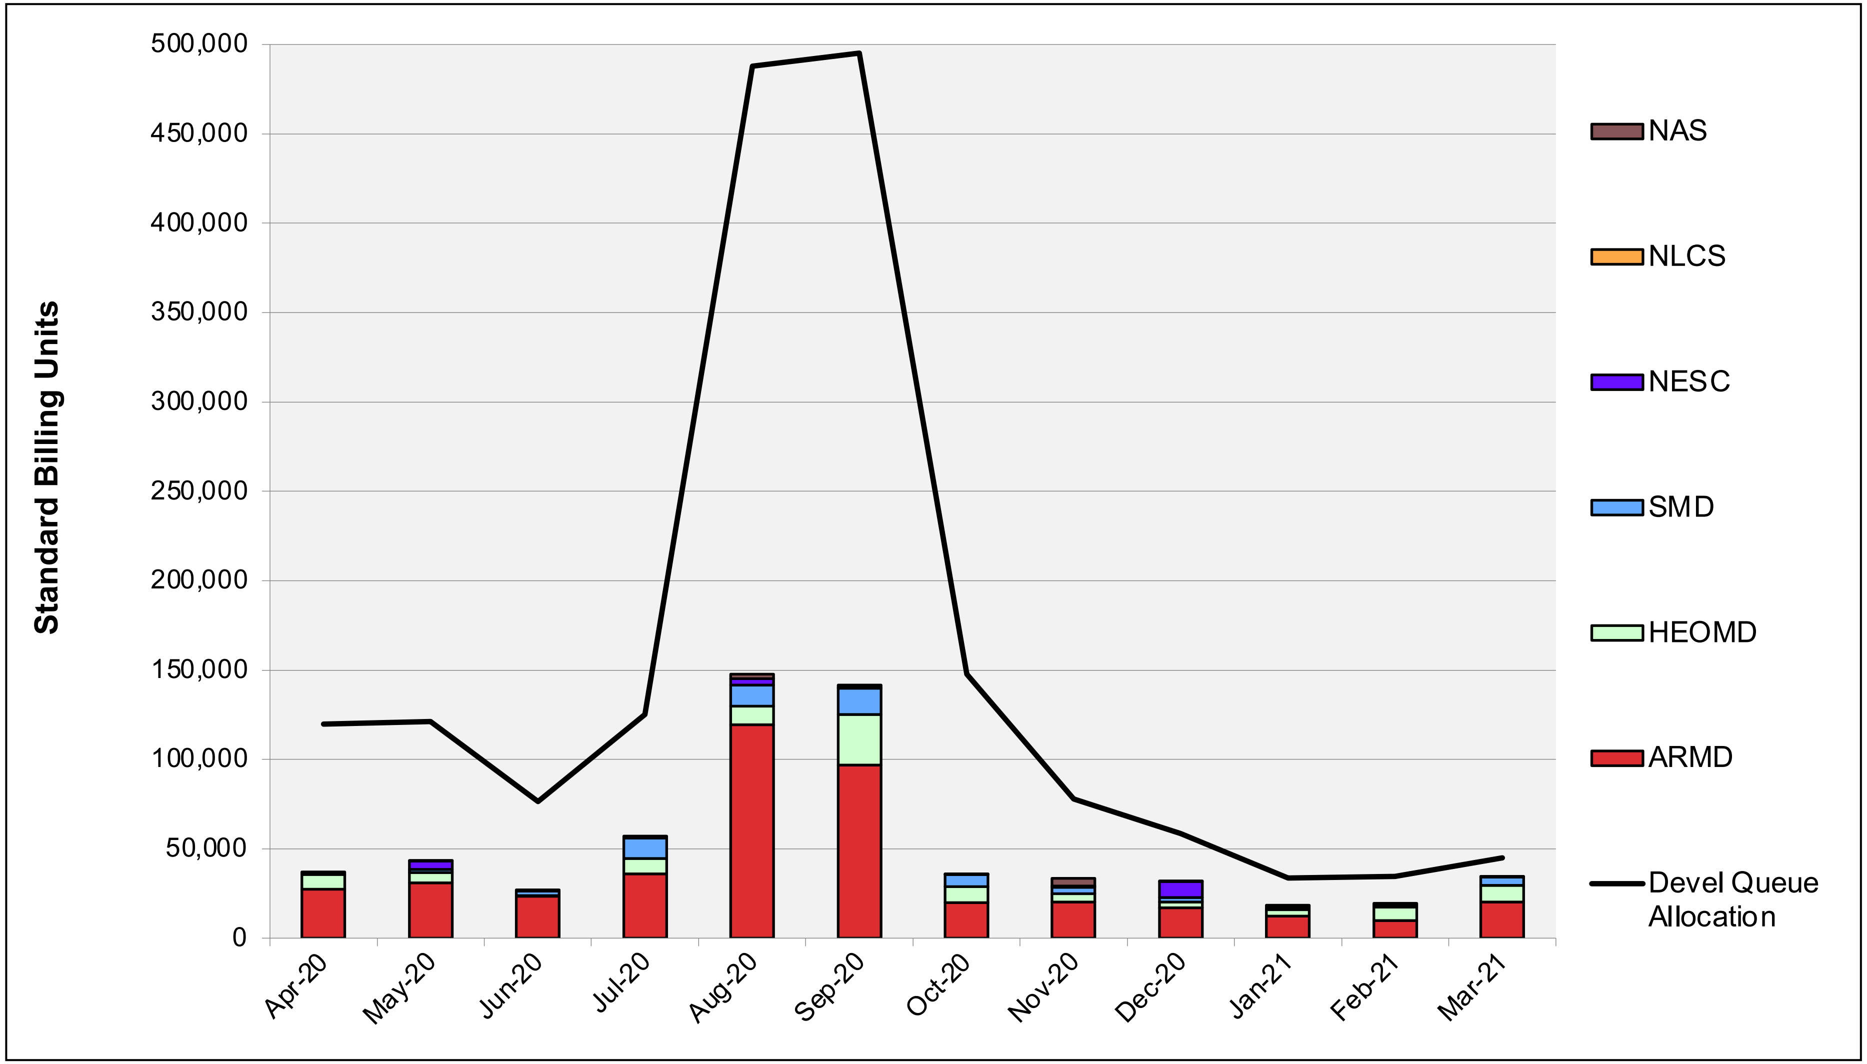

Electra: Monthly Utilization by Job Length

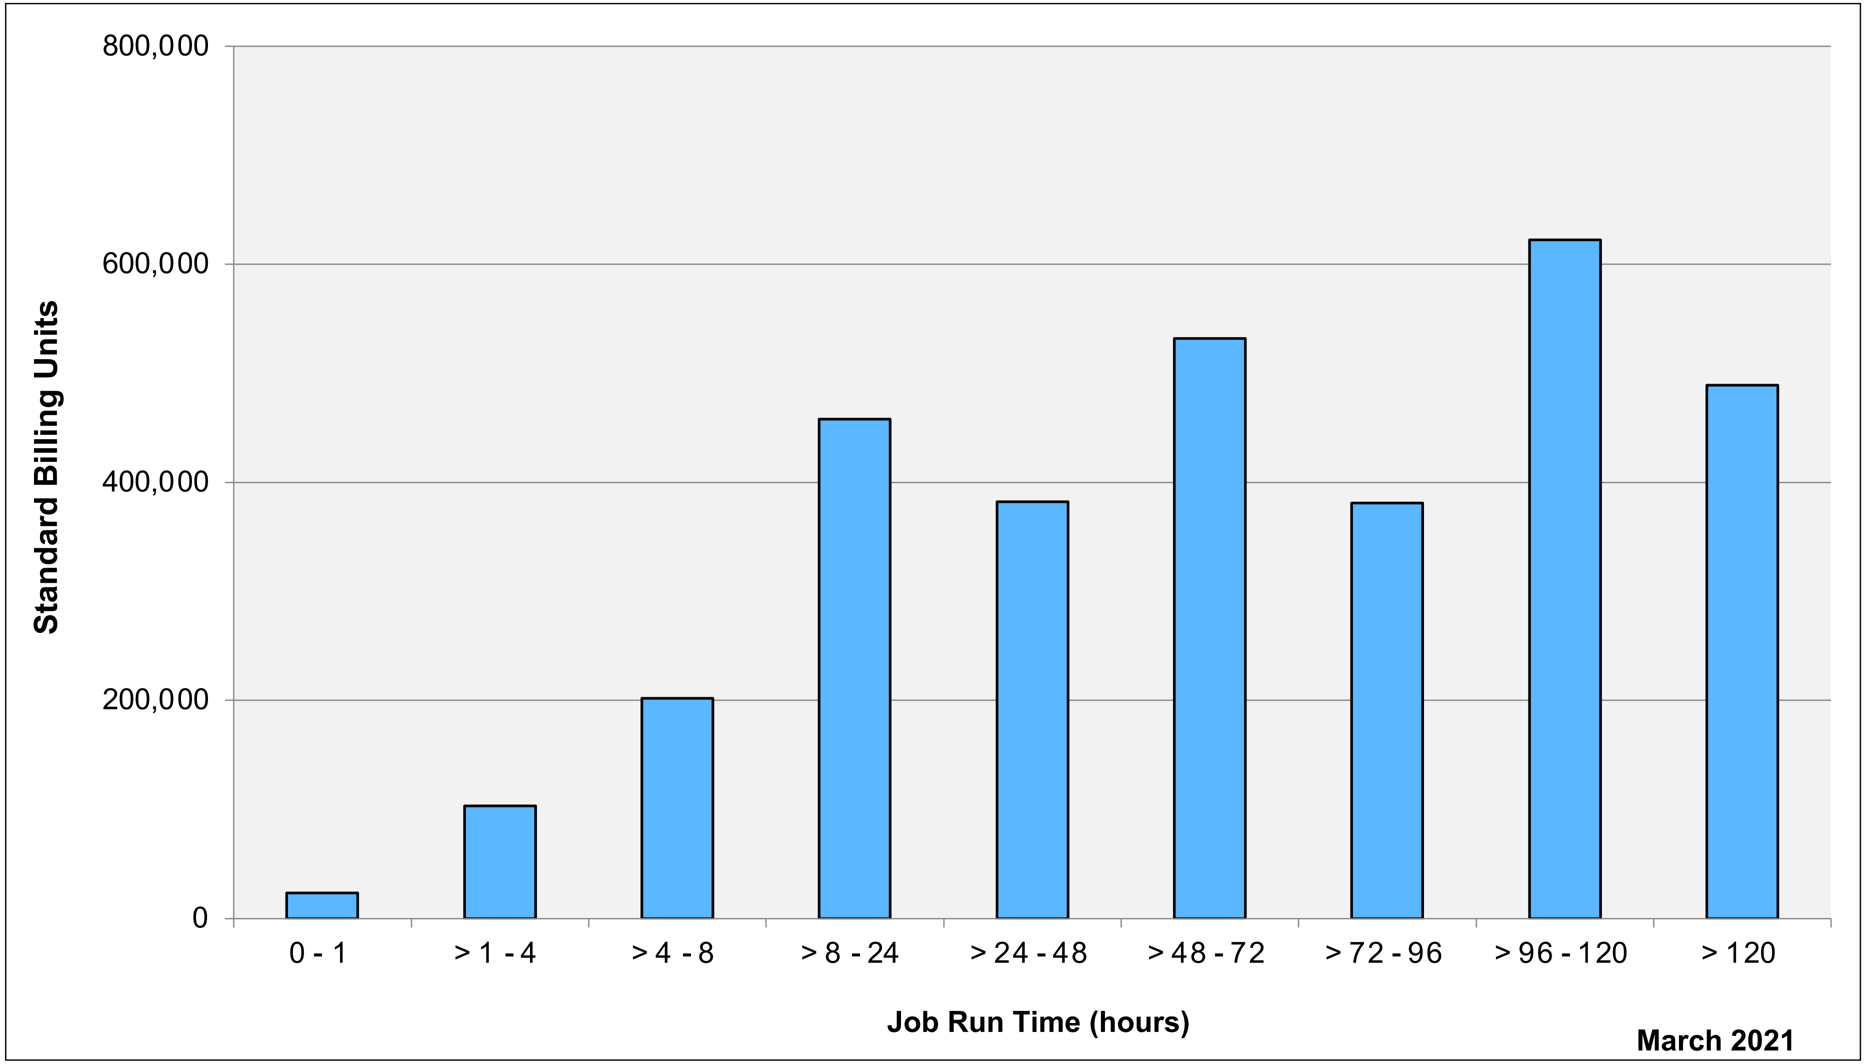

Electra: Monthly Utilization by Job Size

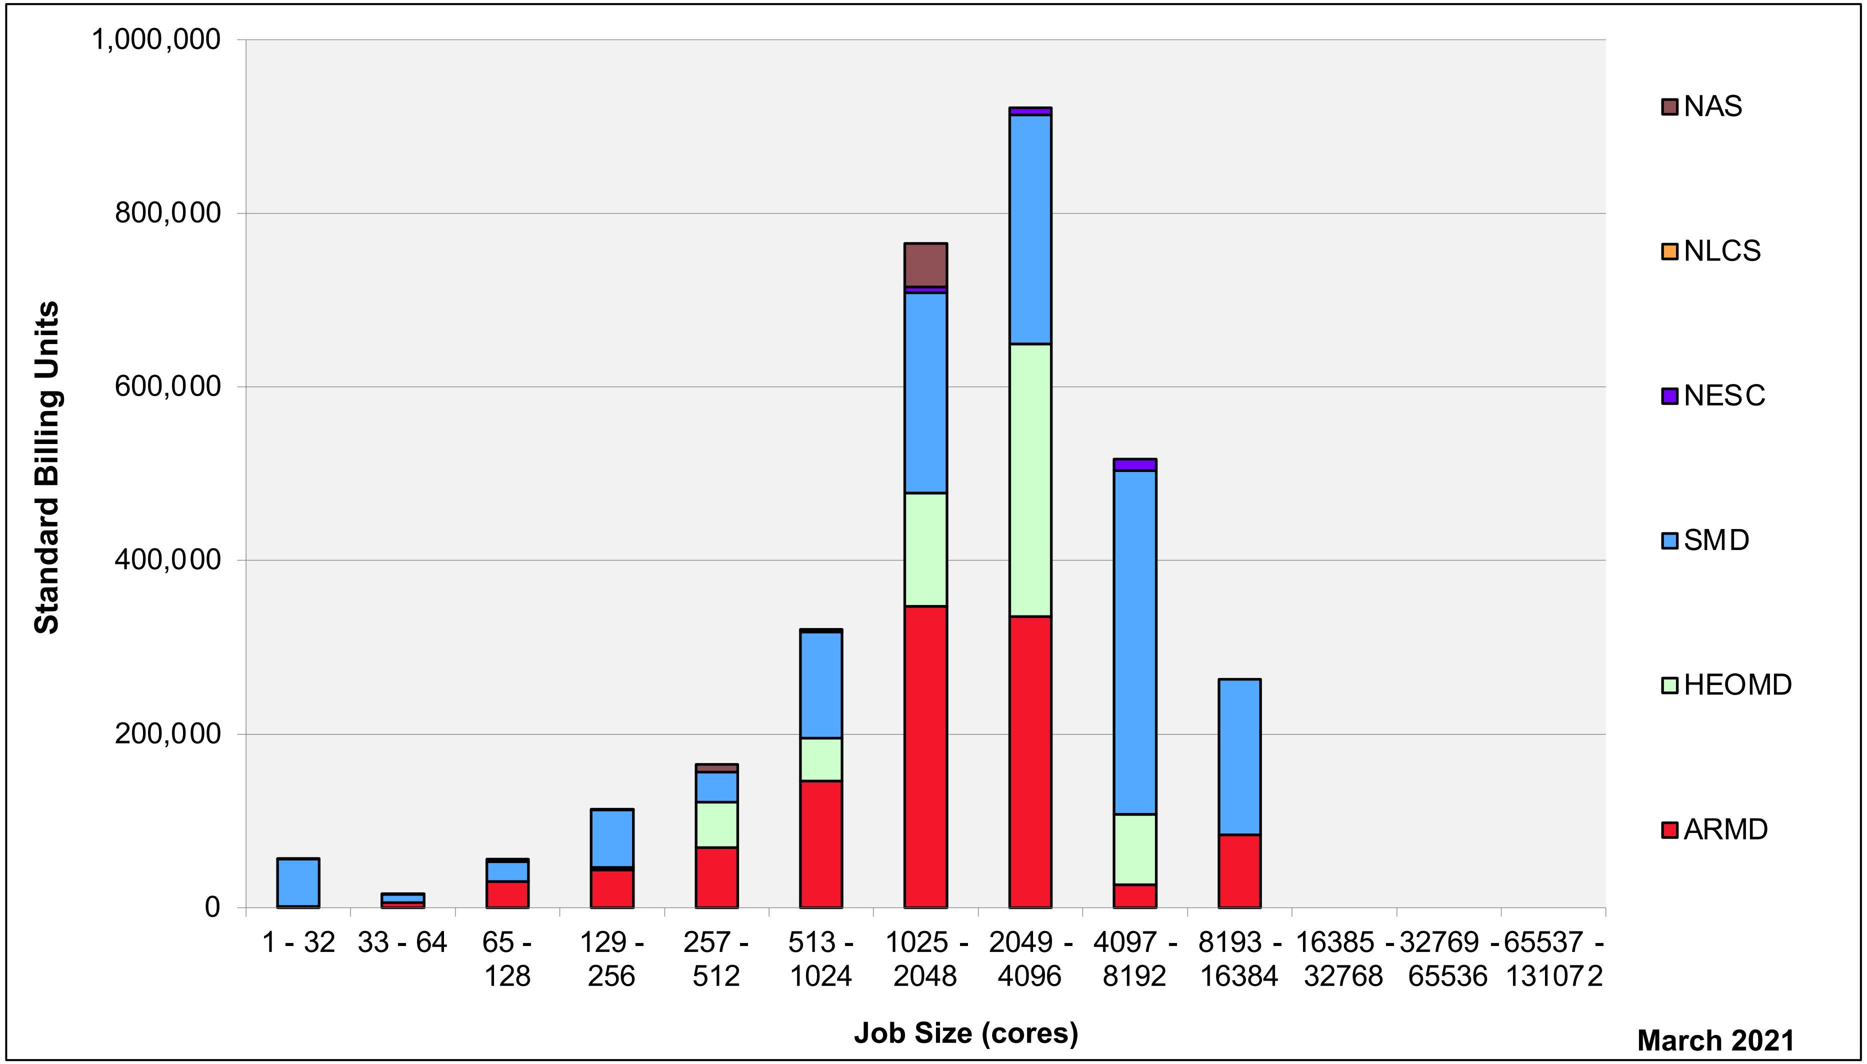

Electra: Monthly Utilization by Size and Length

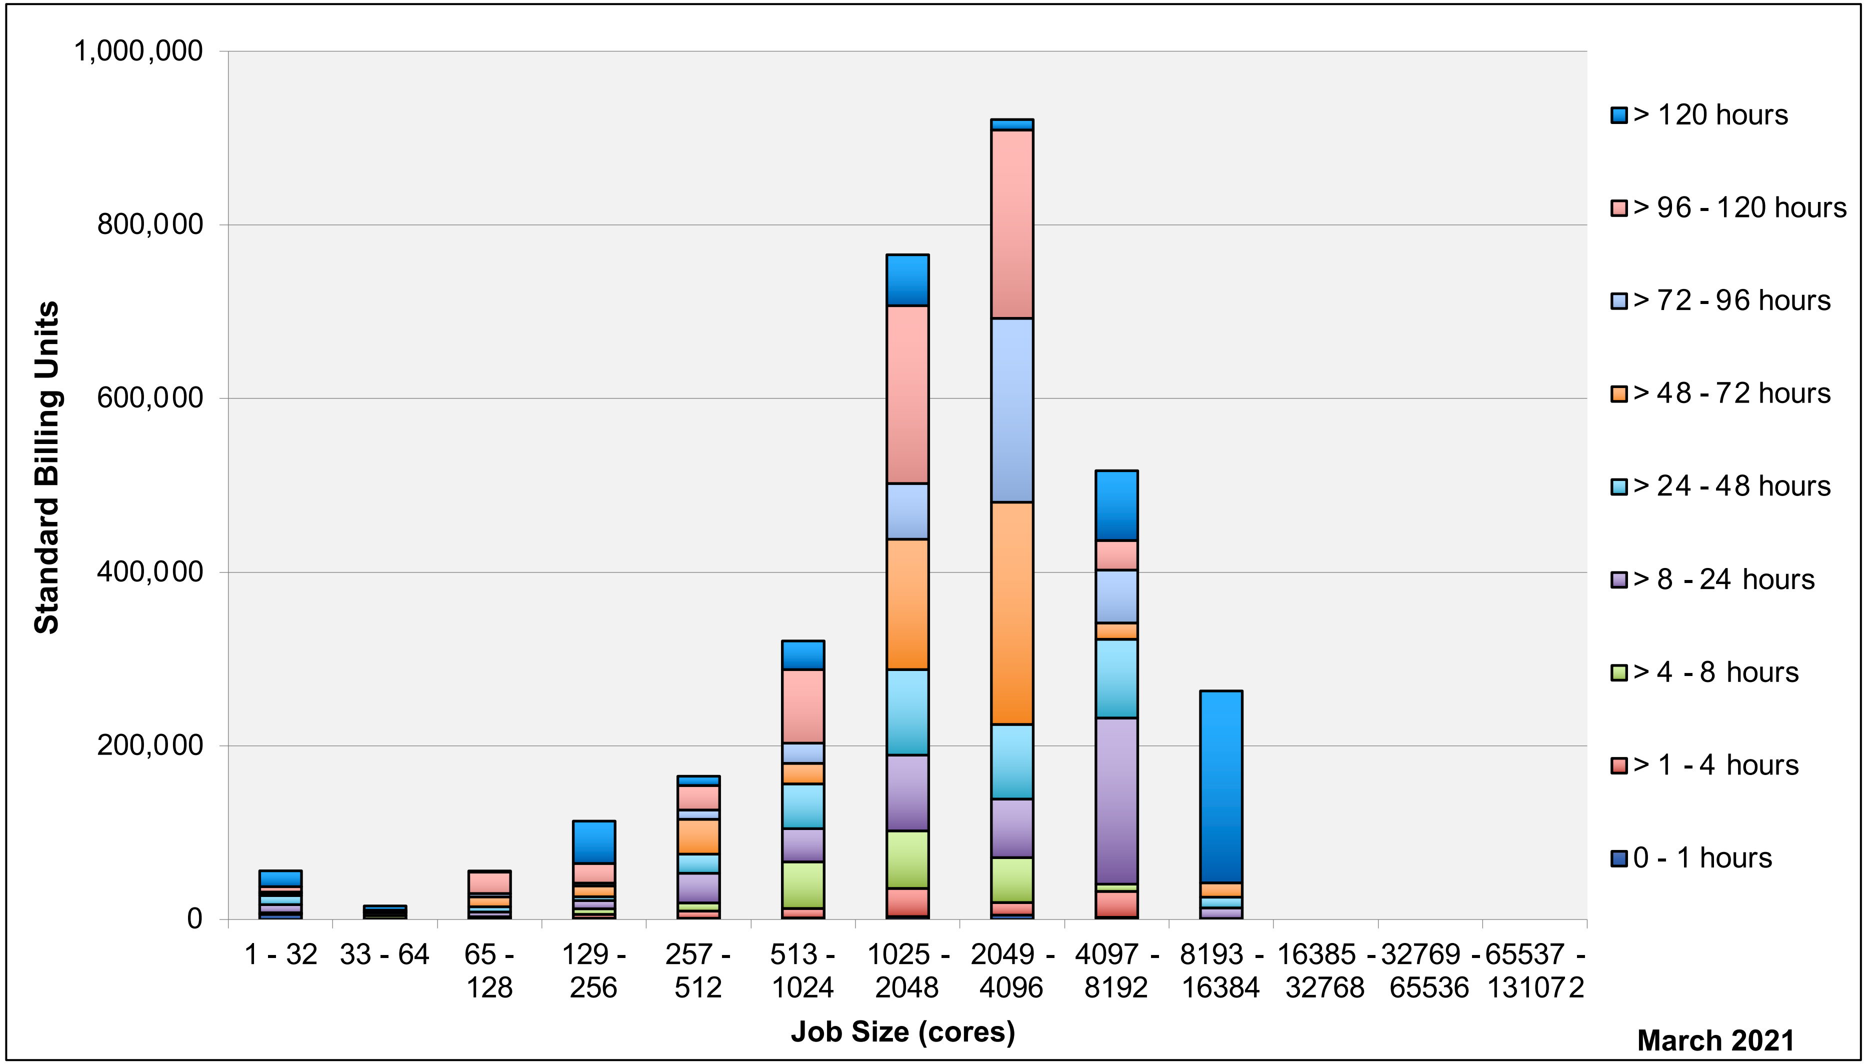

Electra: Average Time to Clear All Jobs

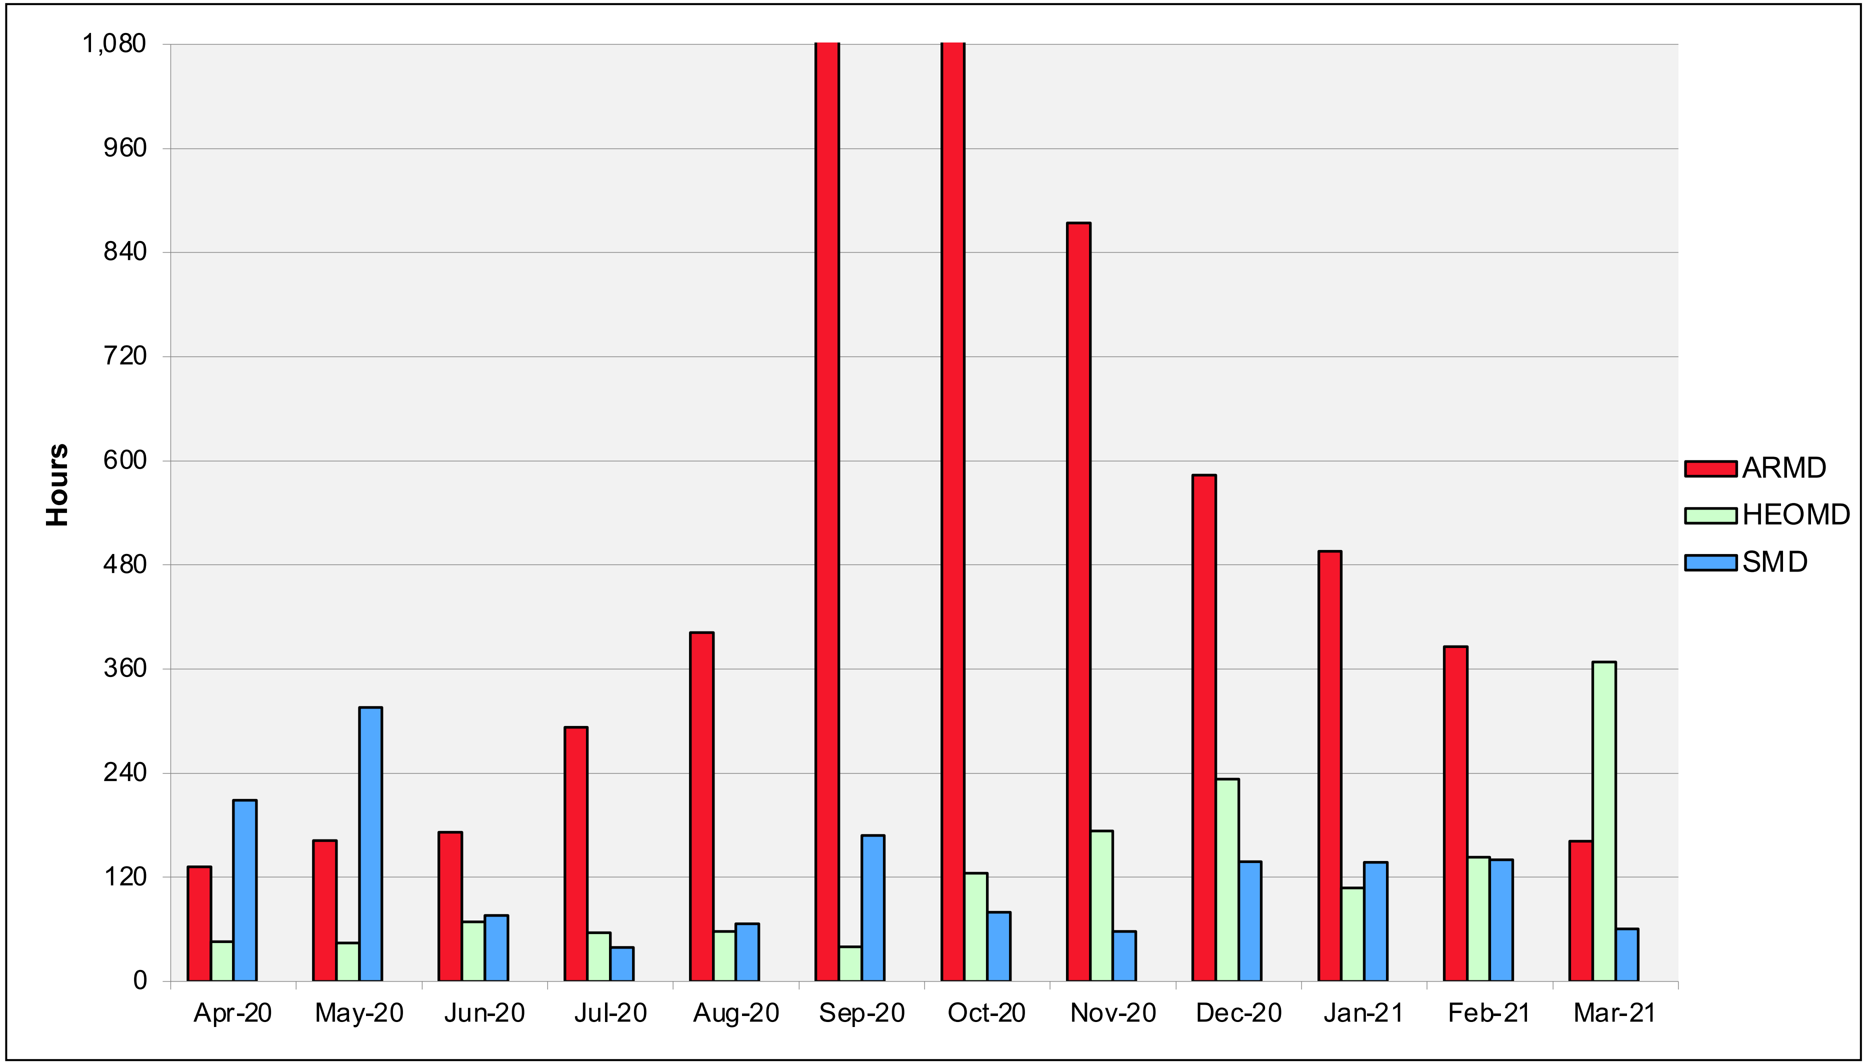

Electra: Average Expansion Factor

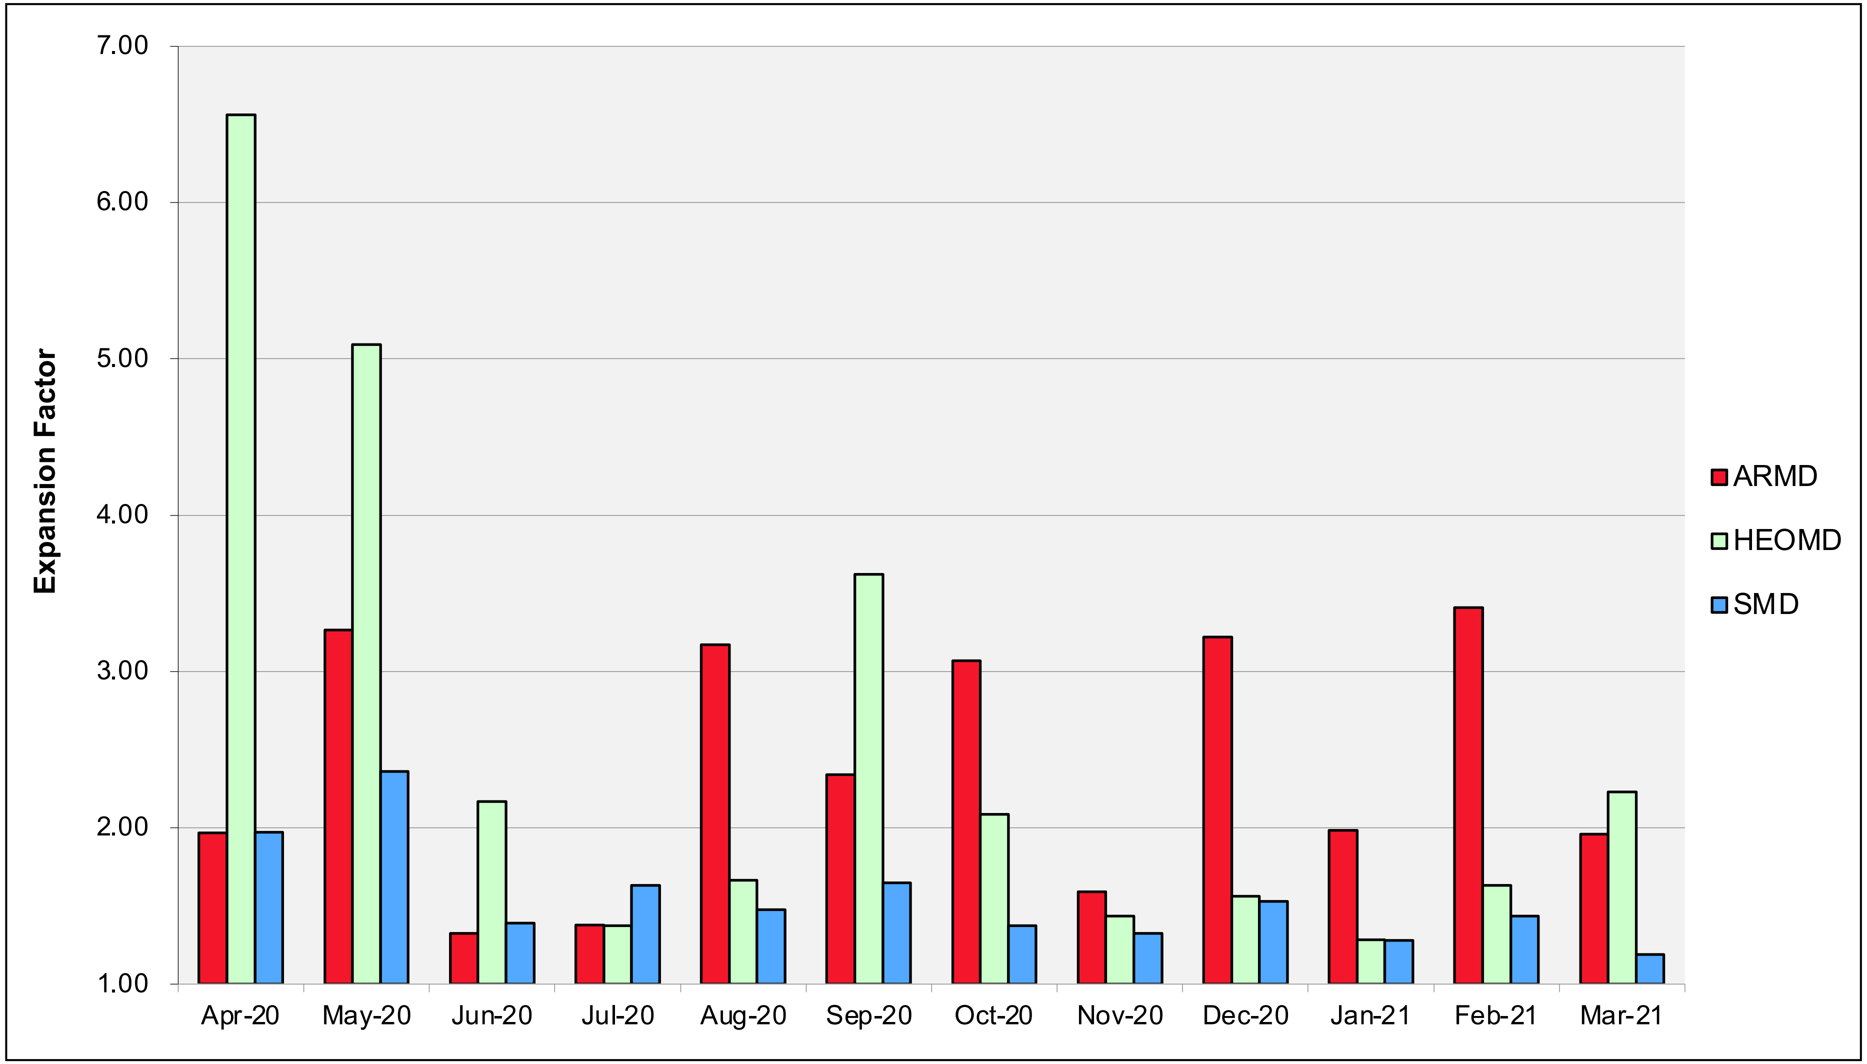

Merope: SBUs Reported, Normalized to 30-Day Month

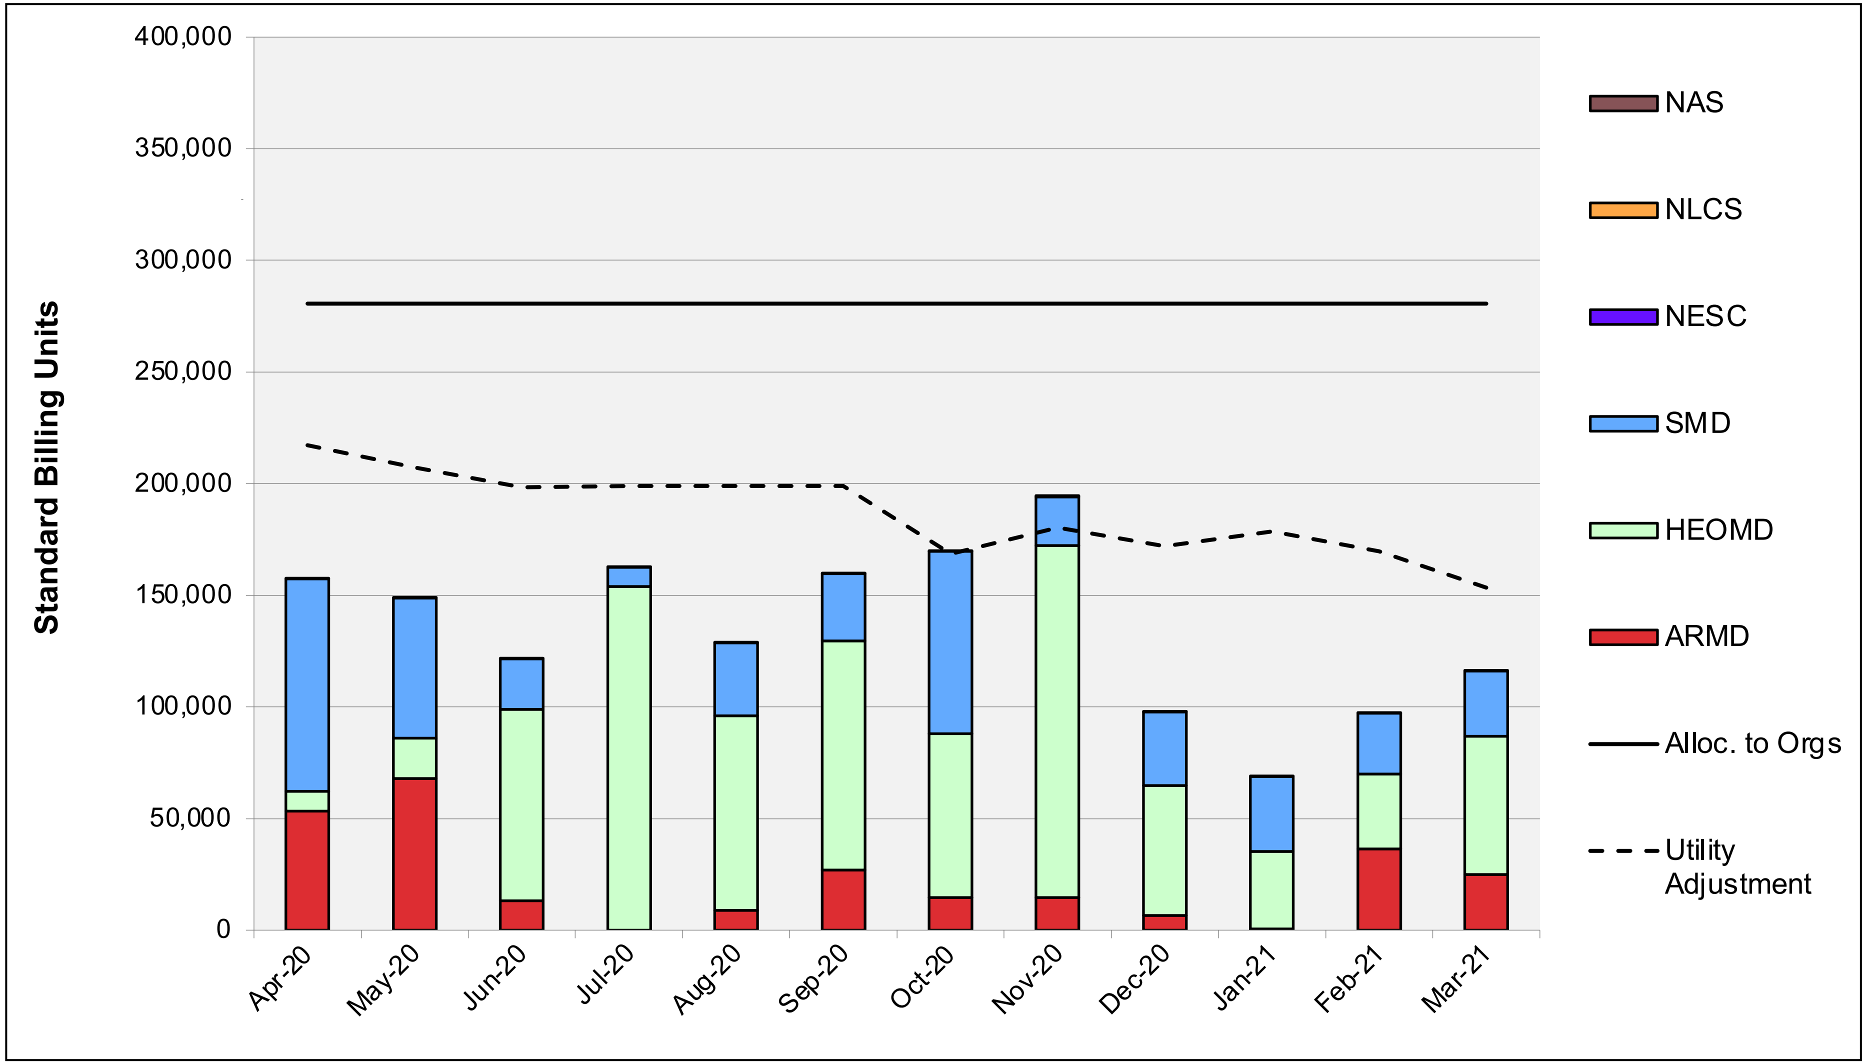

Merope: Monthly Utilization by Job Length

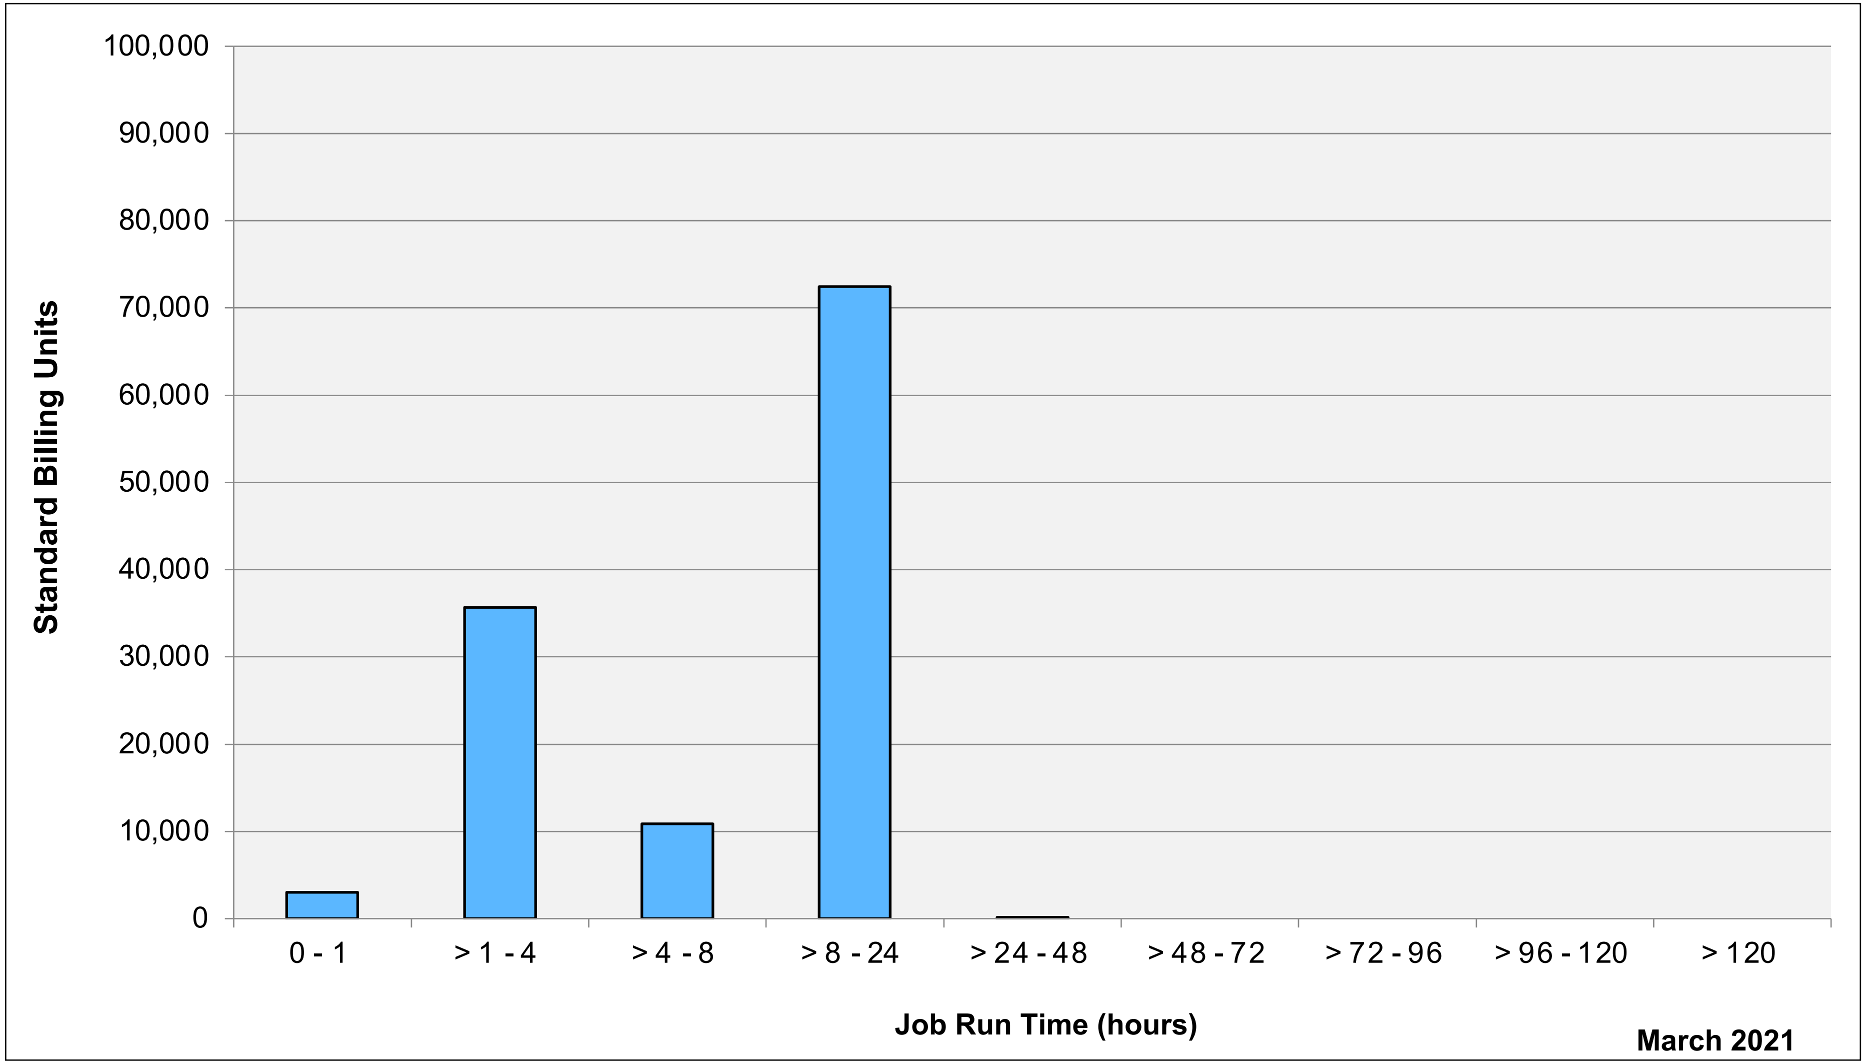

Merope: Monthly Utilization by Job Size

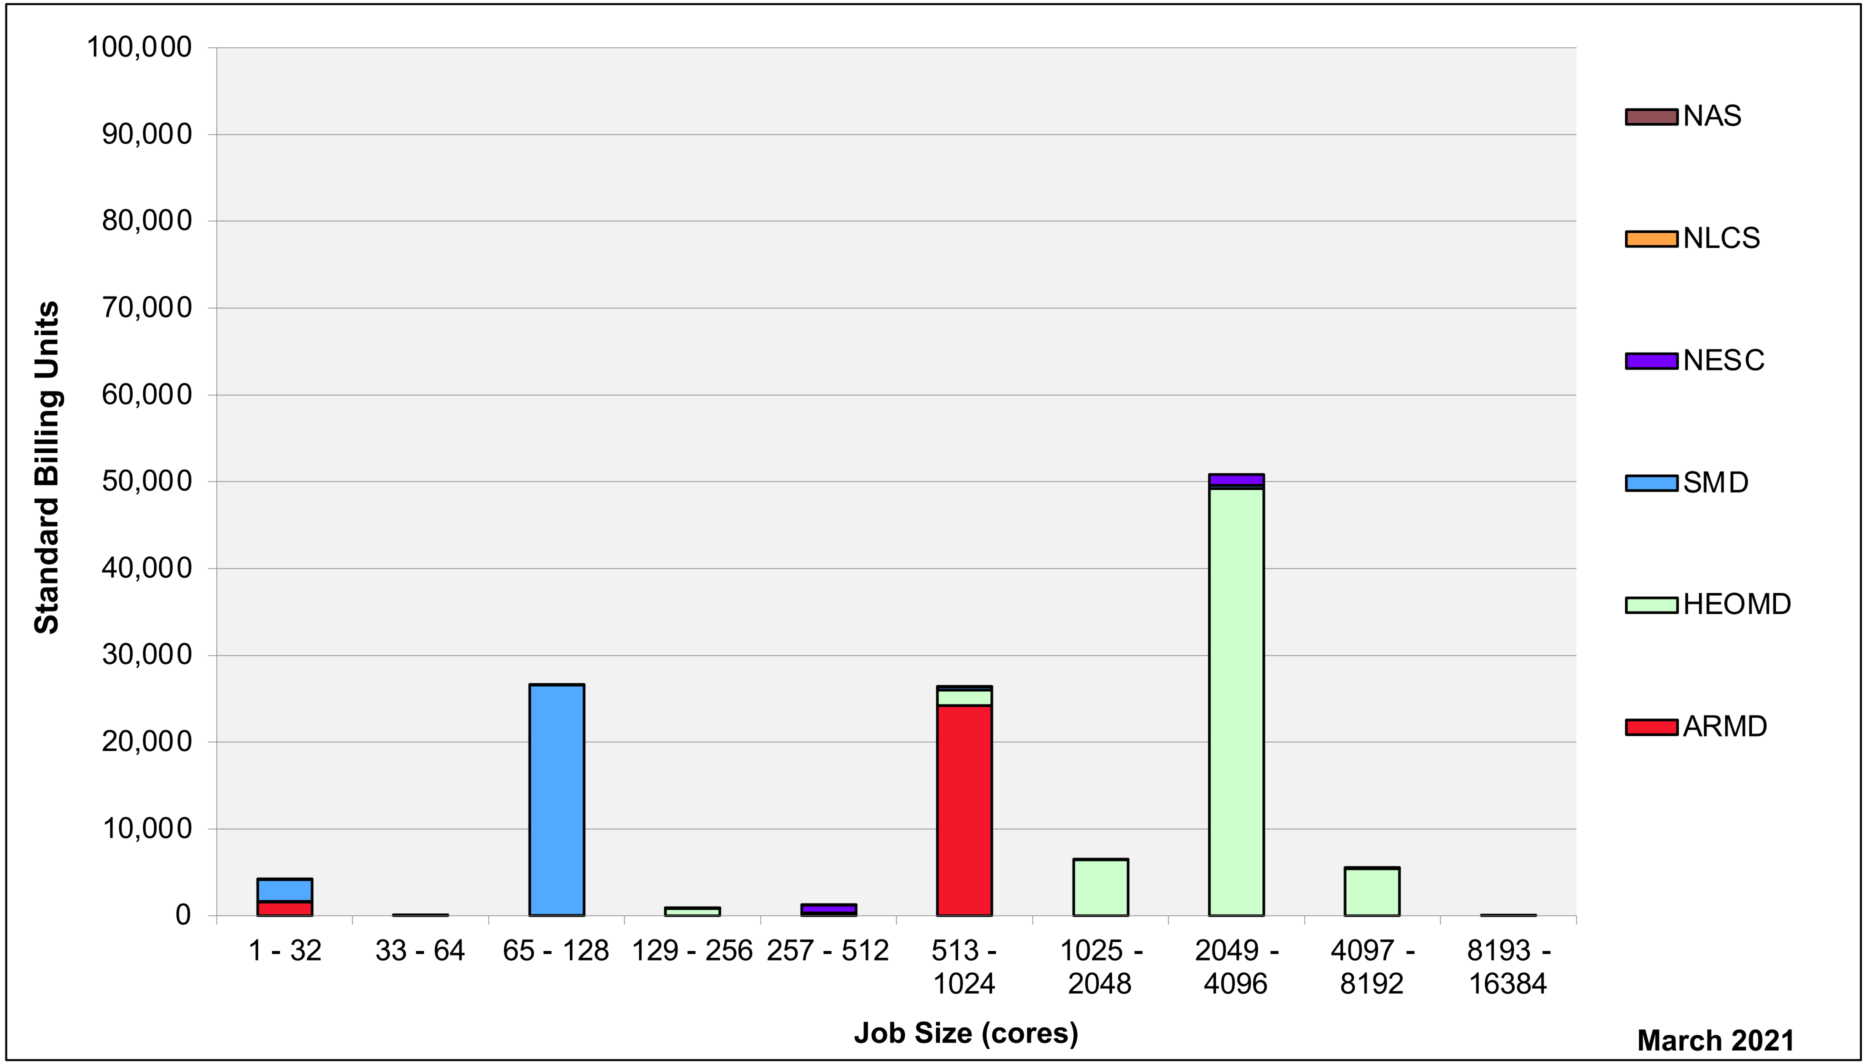

Merope: Monthly Utilization by Size and Length

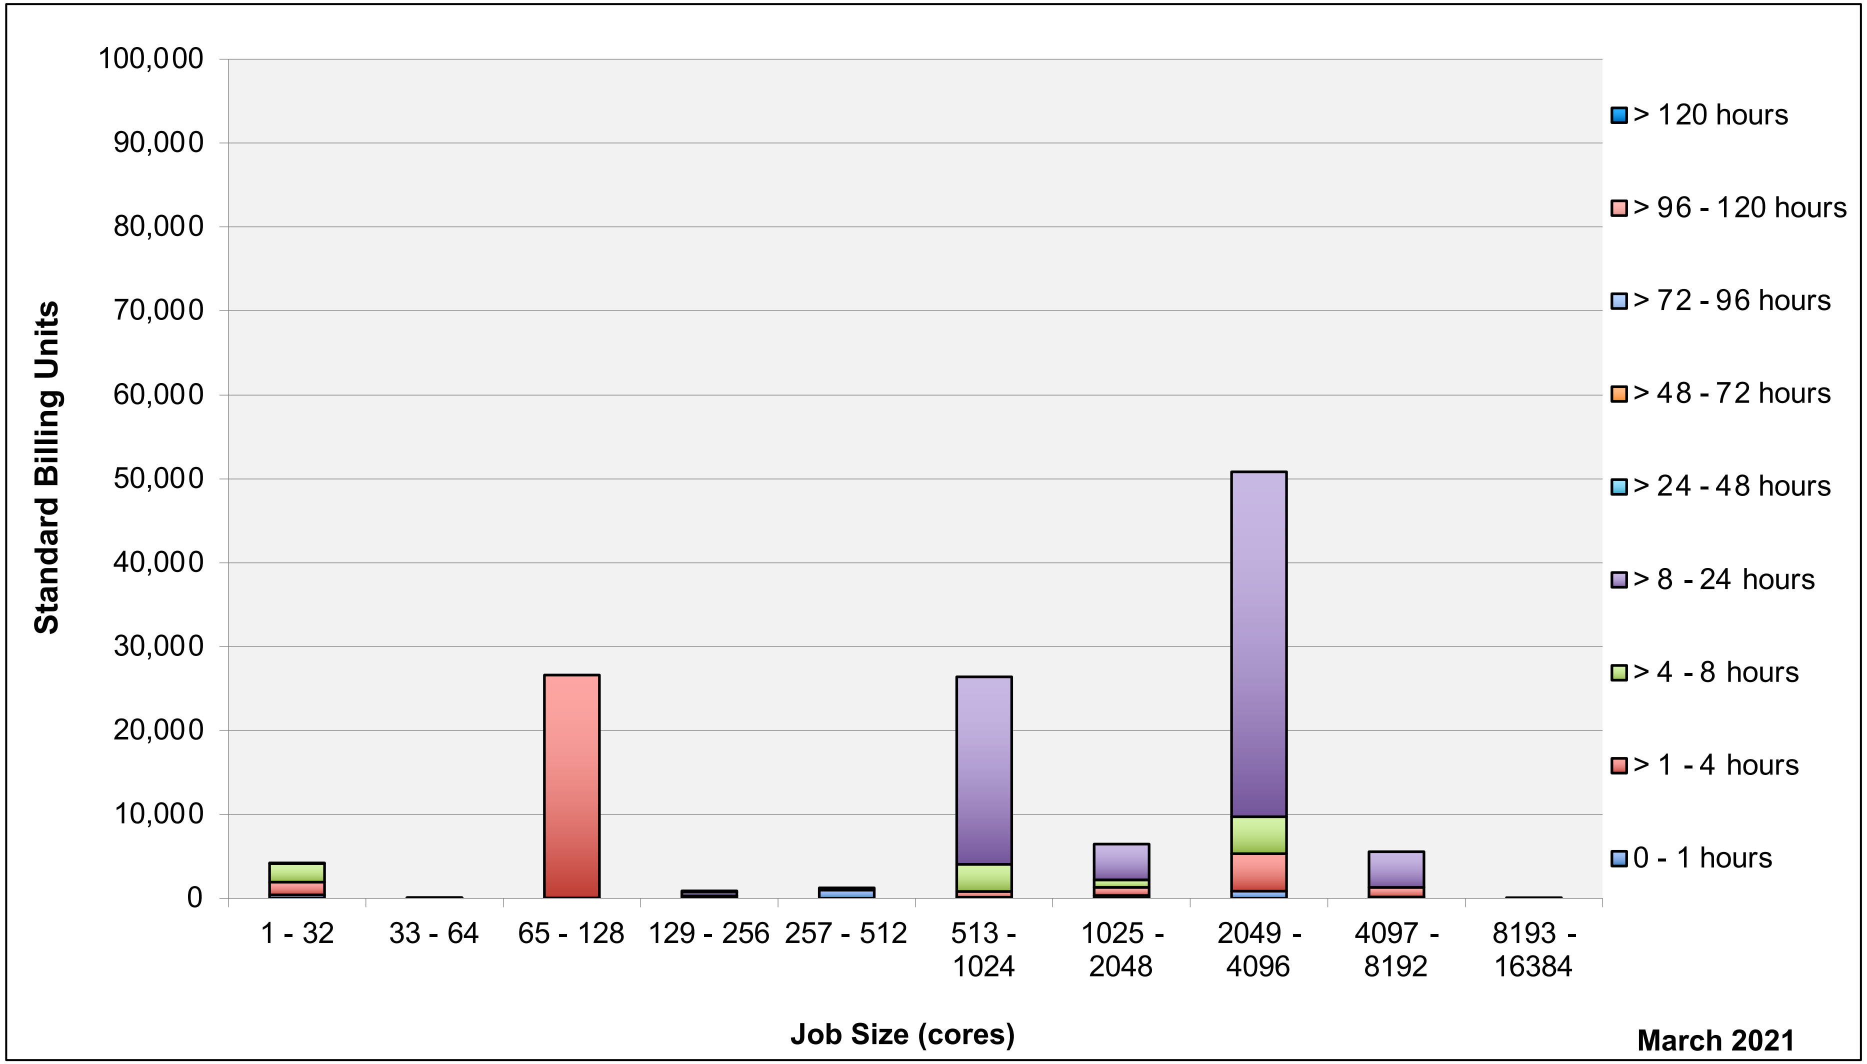

Merope: Average Time to Clear All Jobs

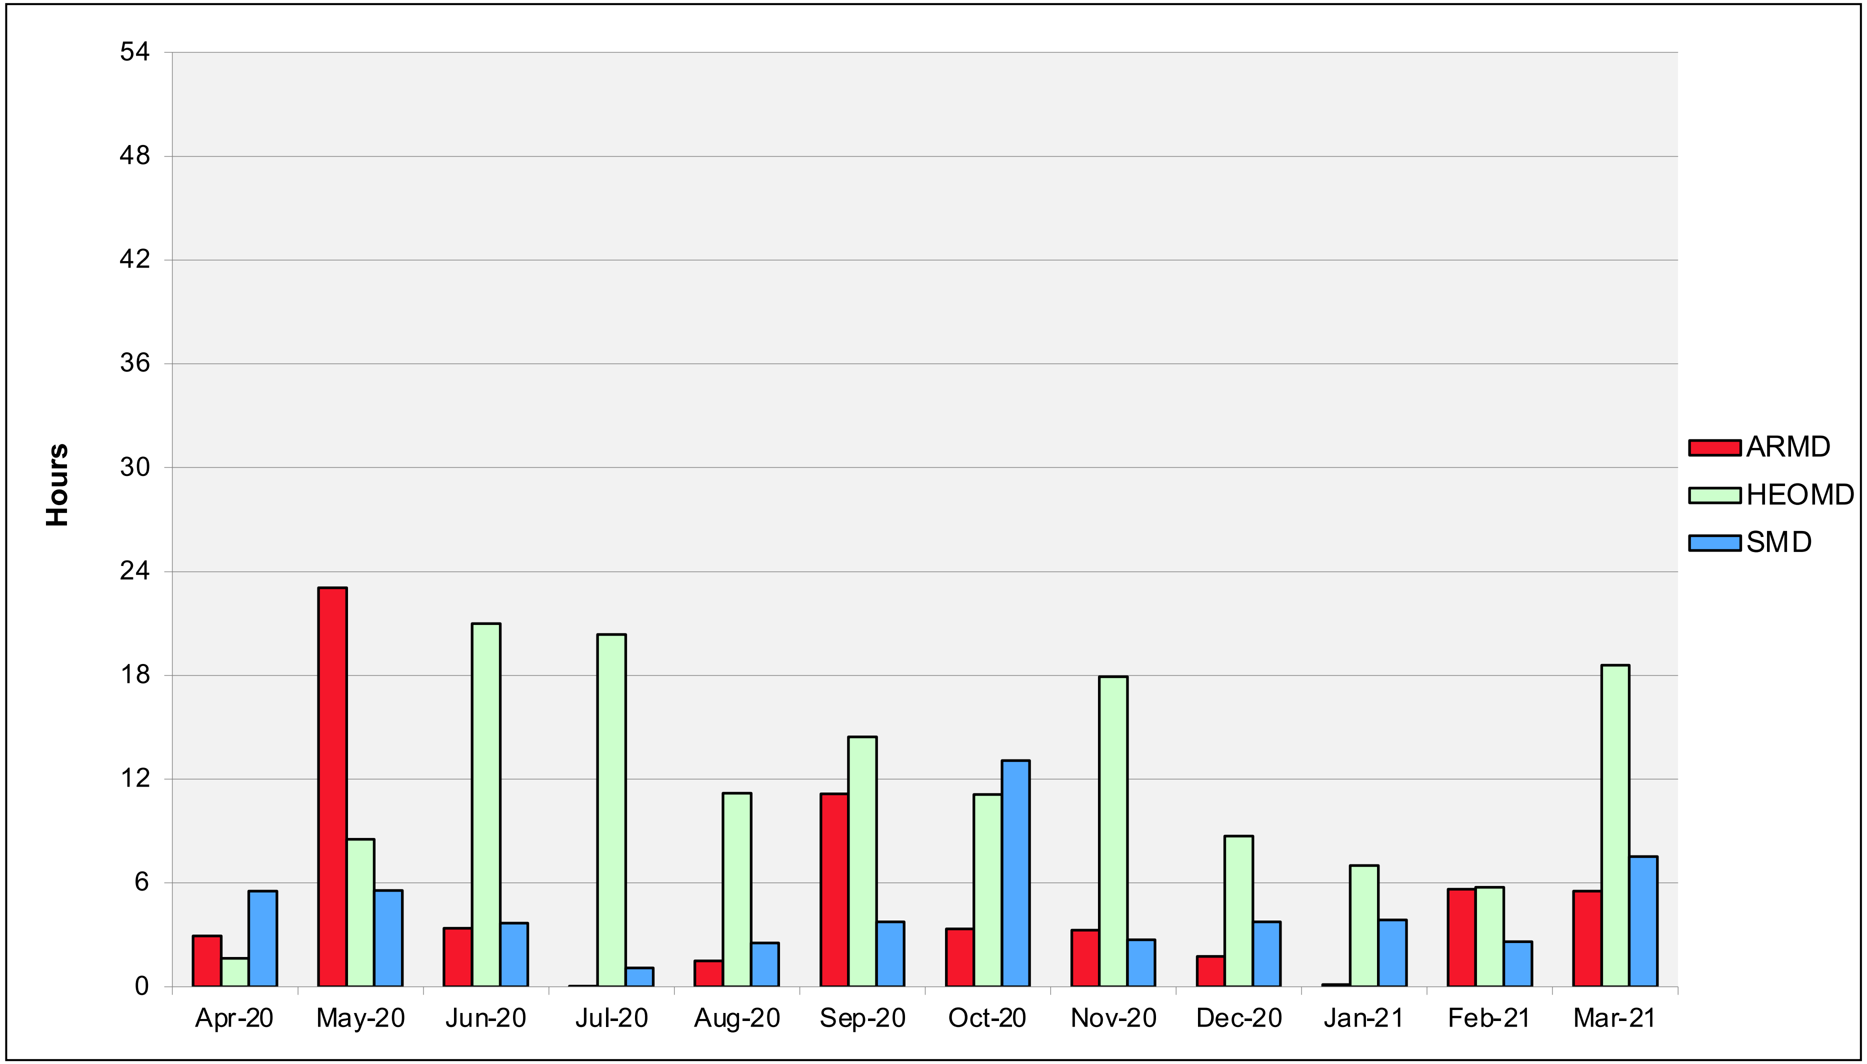

Merope: Average Expansion Factor

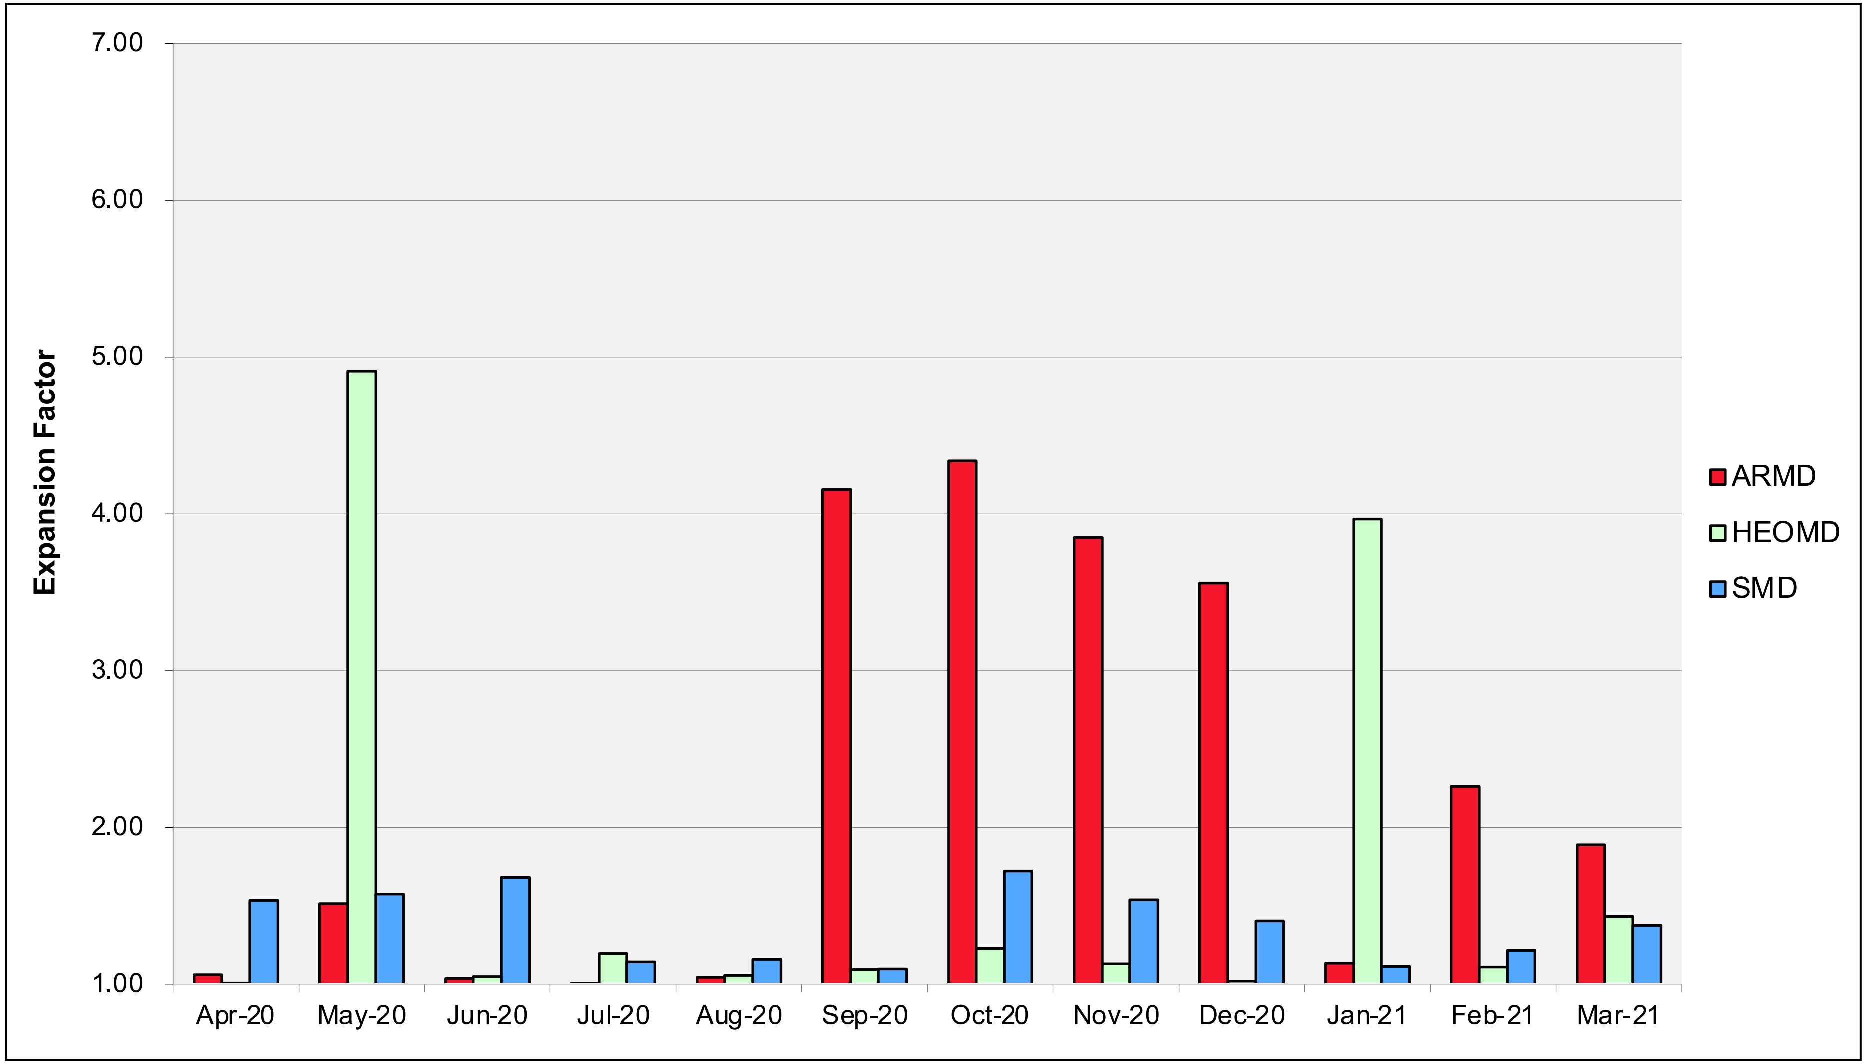

Endeavour: SBUs Reported, Normalized to 30-Day Month

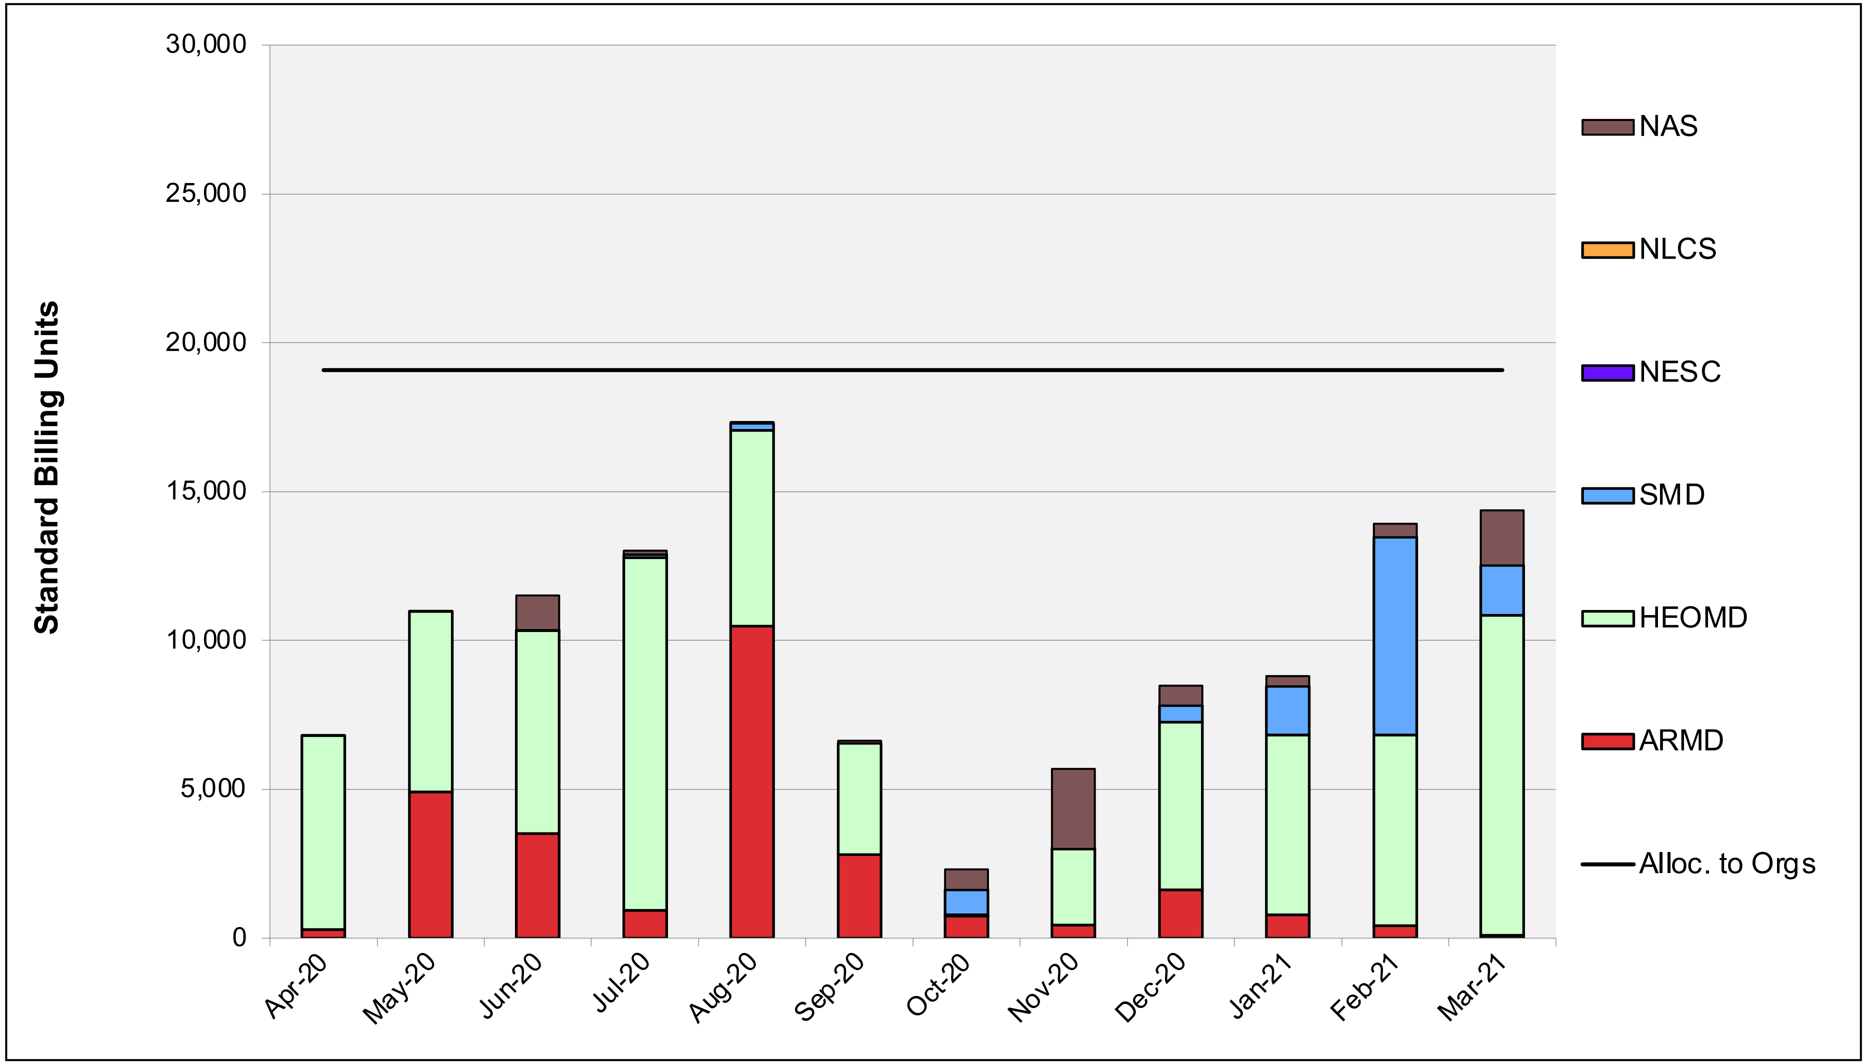

Endeavour: Monthly Utilization by Job Length

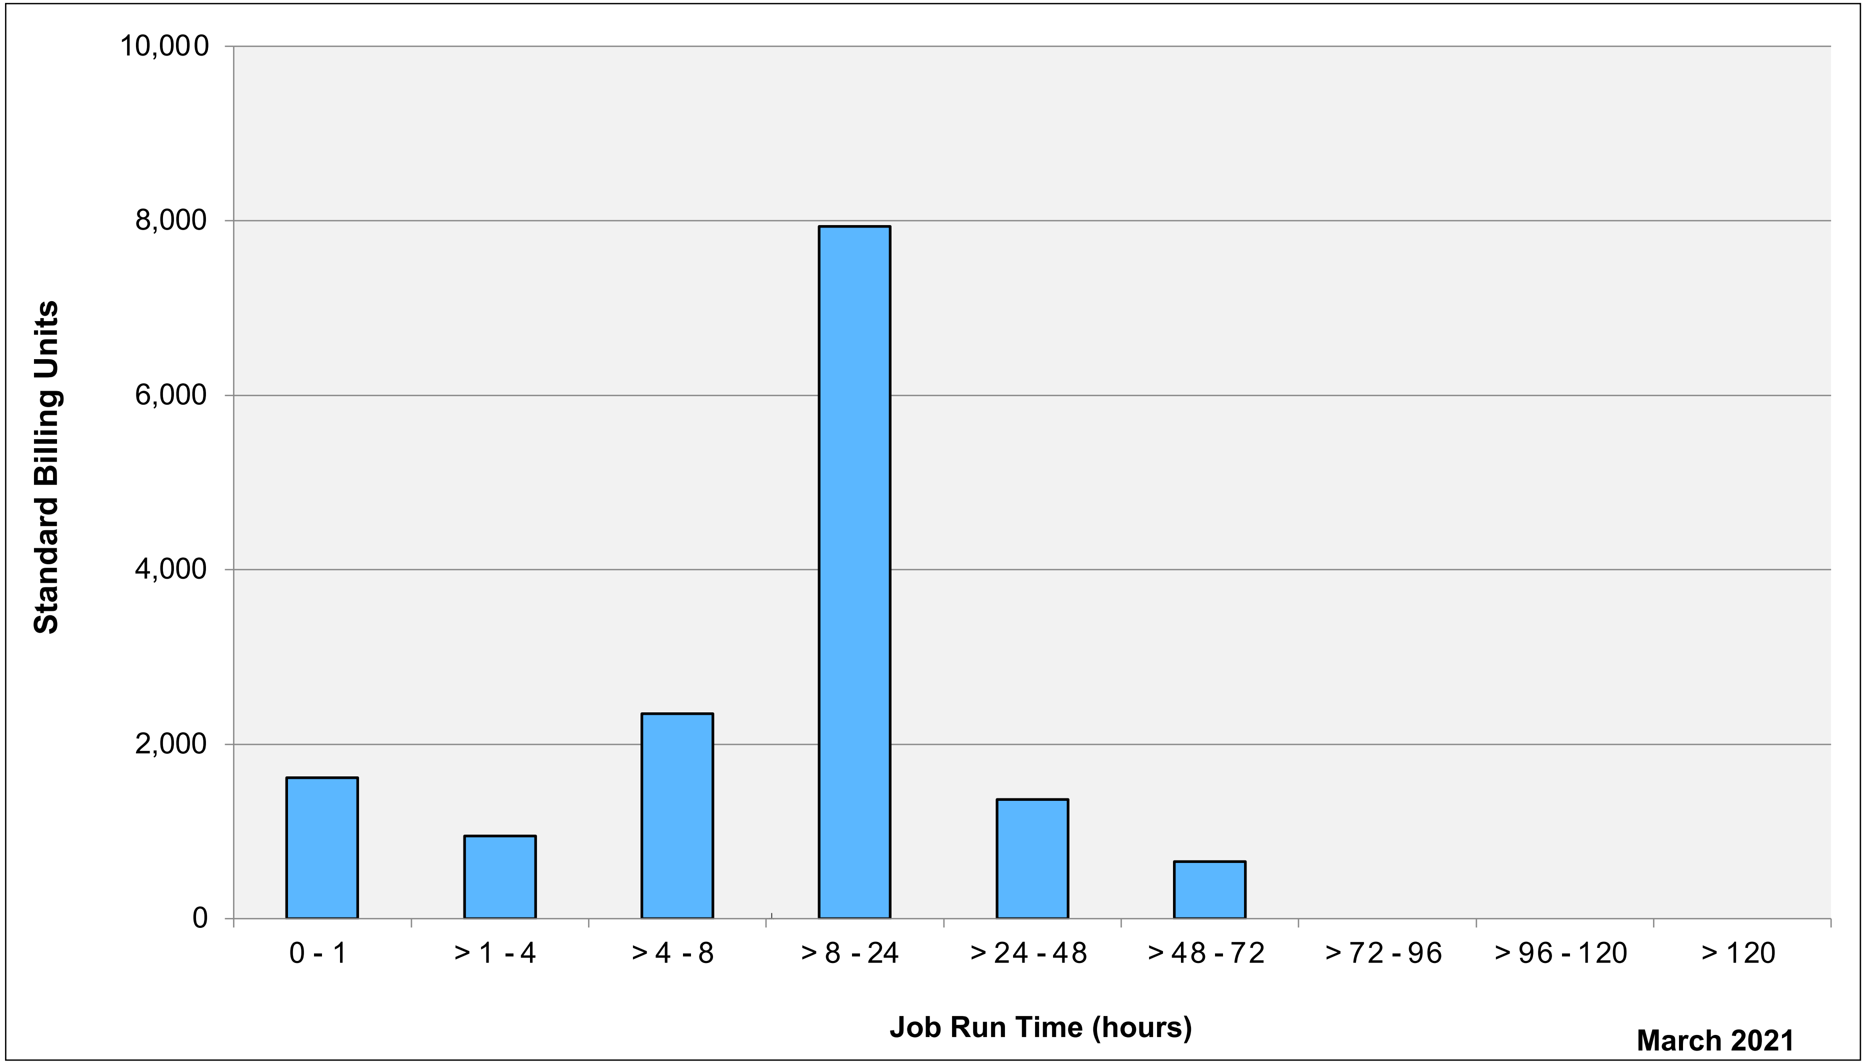

Endeavour: Monthly Utilization by Job Size

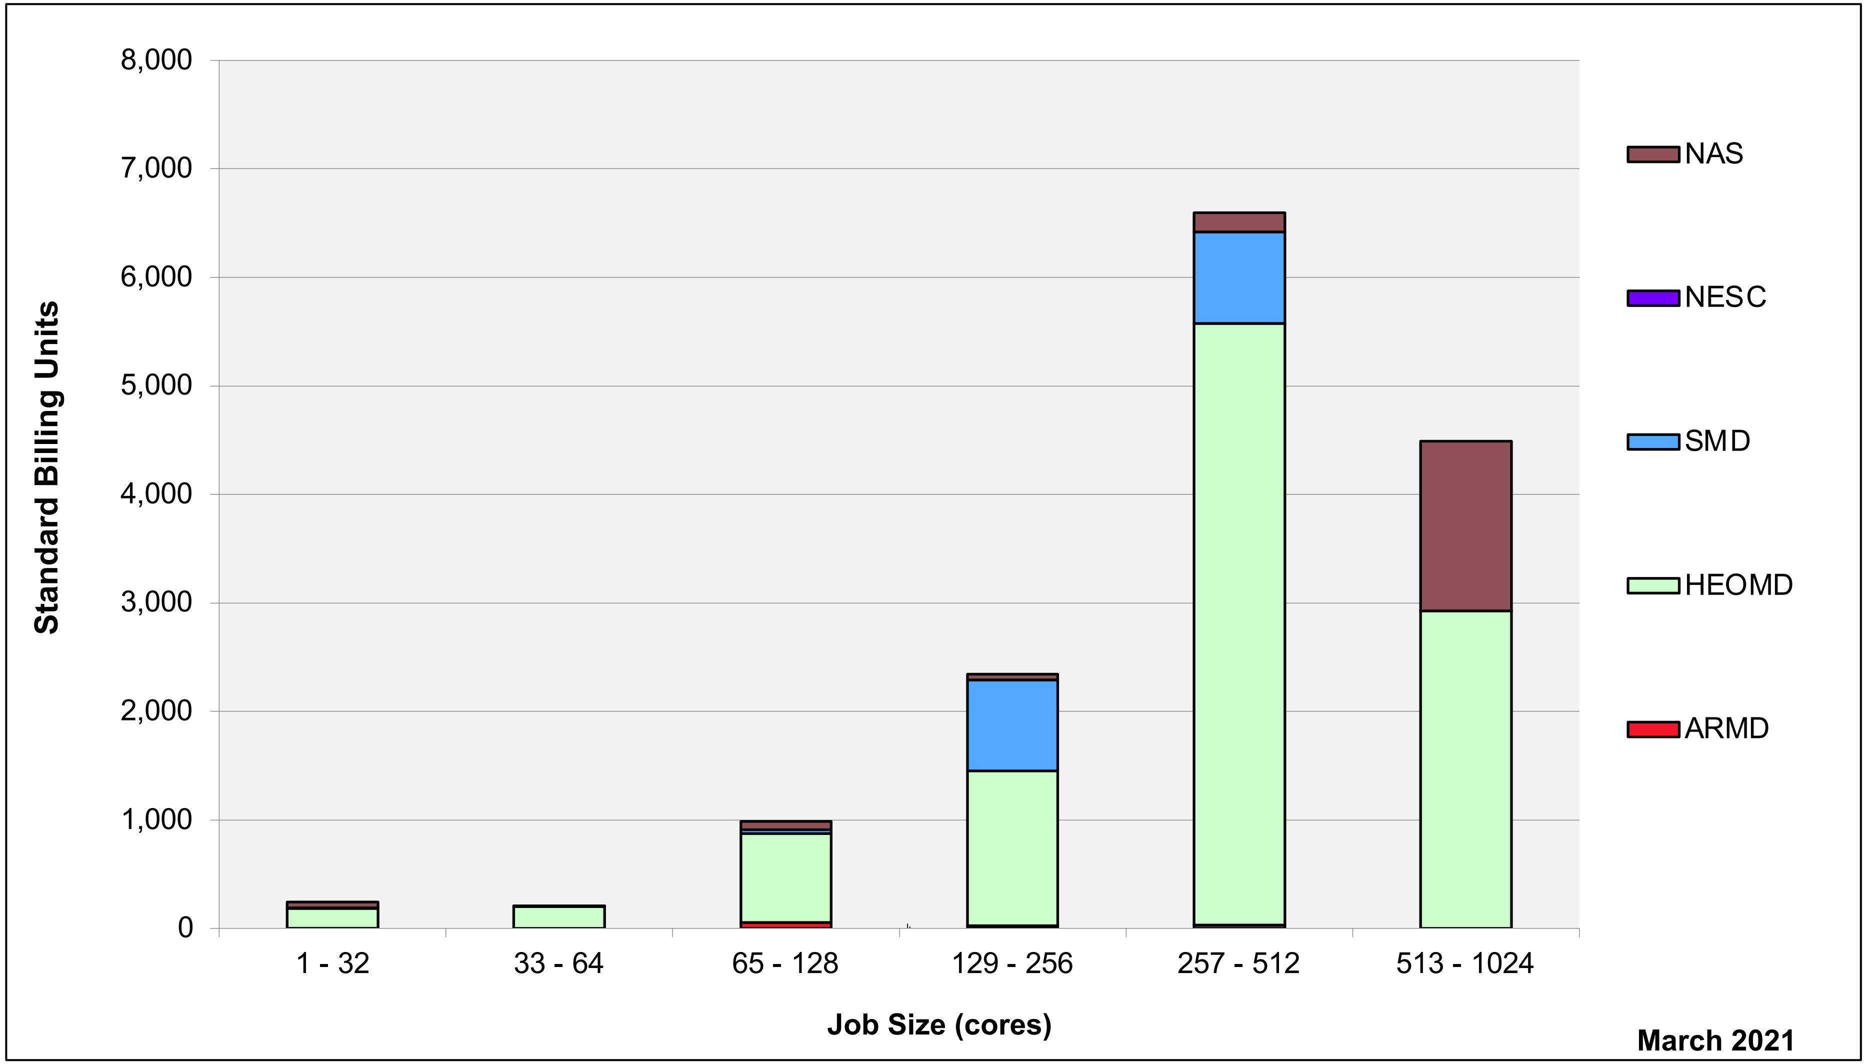

Endeavour: Monthly Utilization by Size and Length

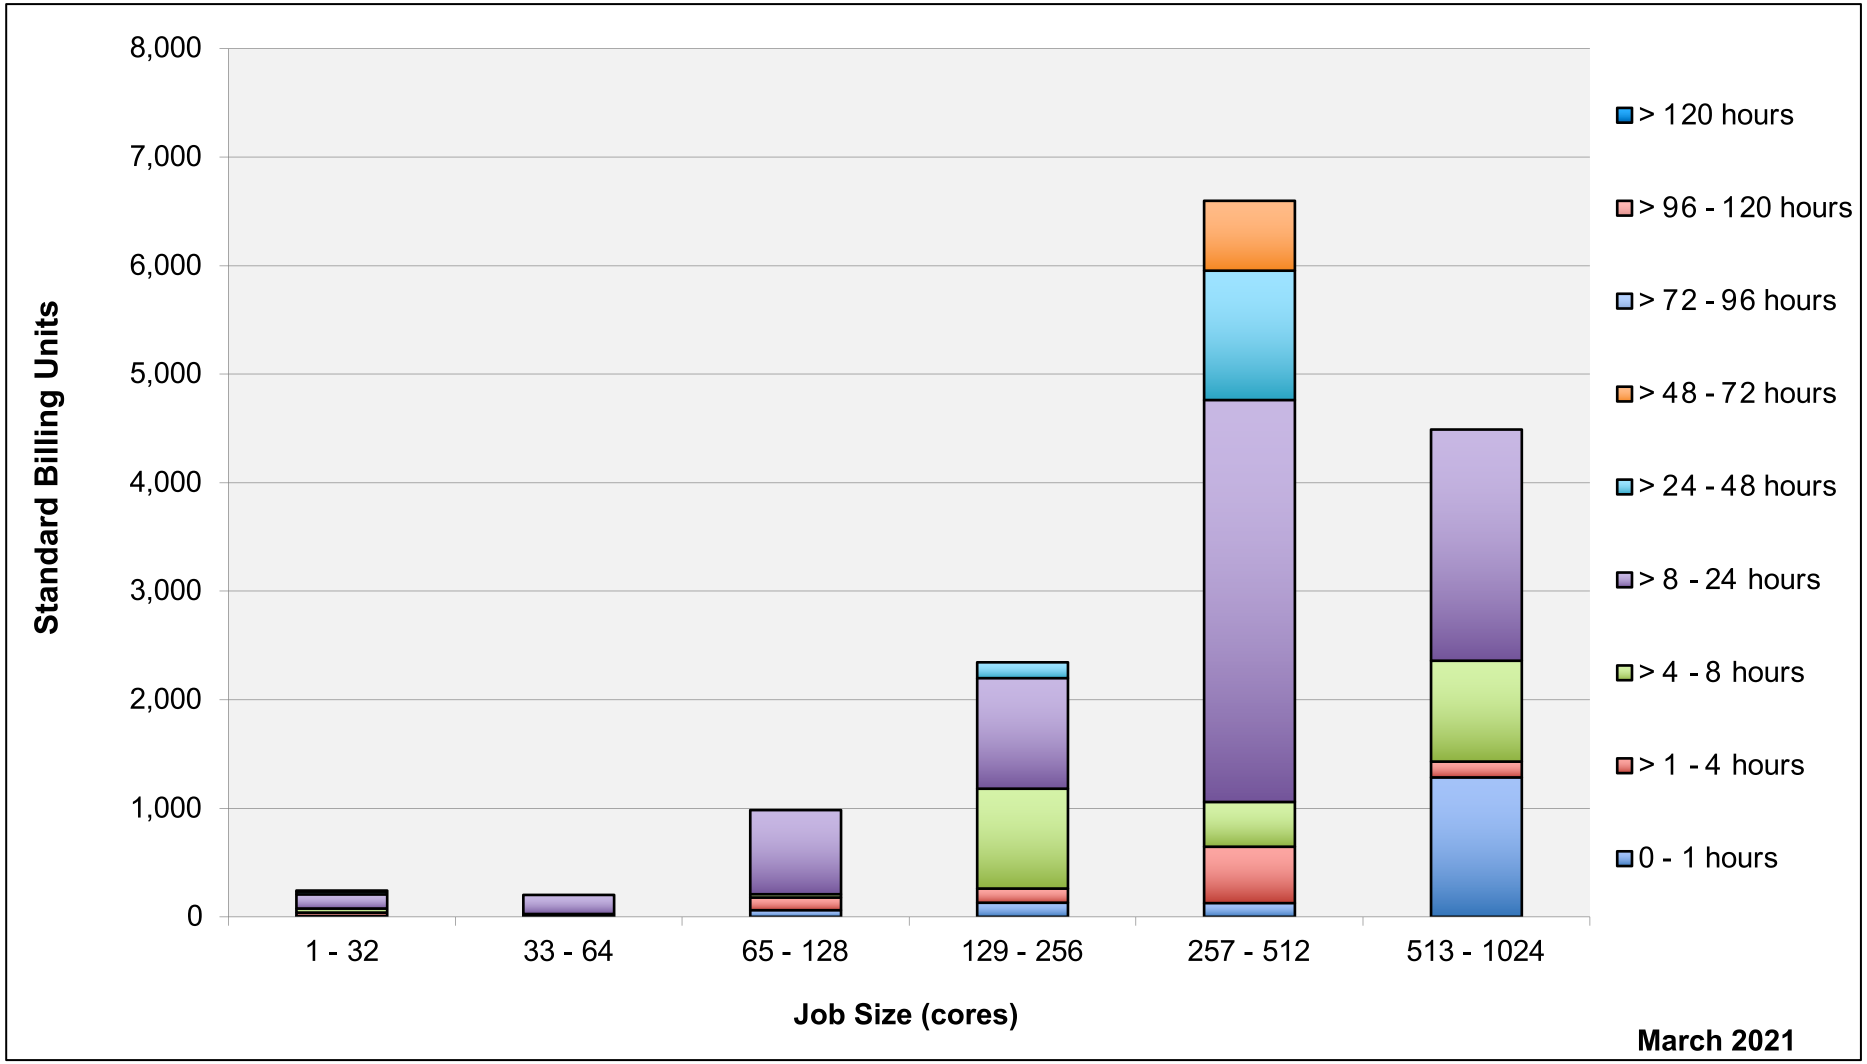

Endeavour: Average Time to Clear All Jobs

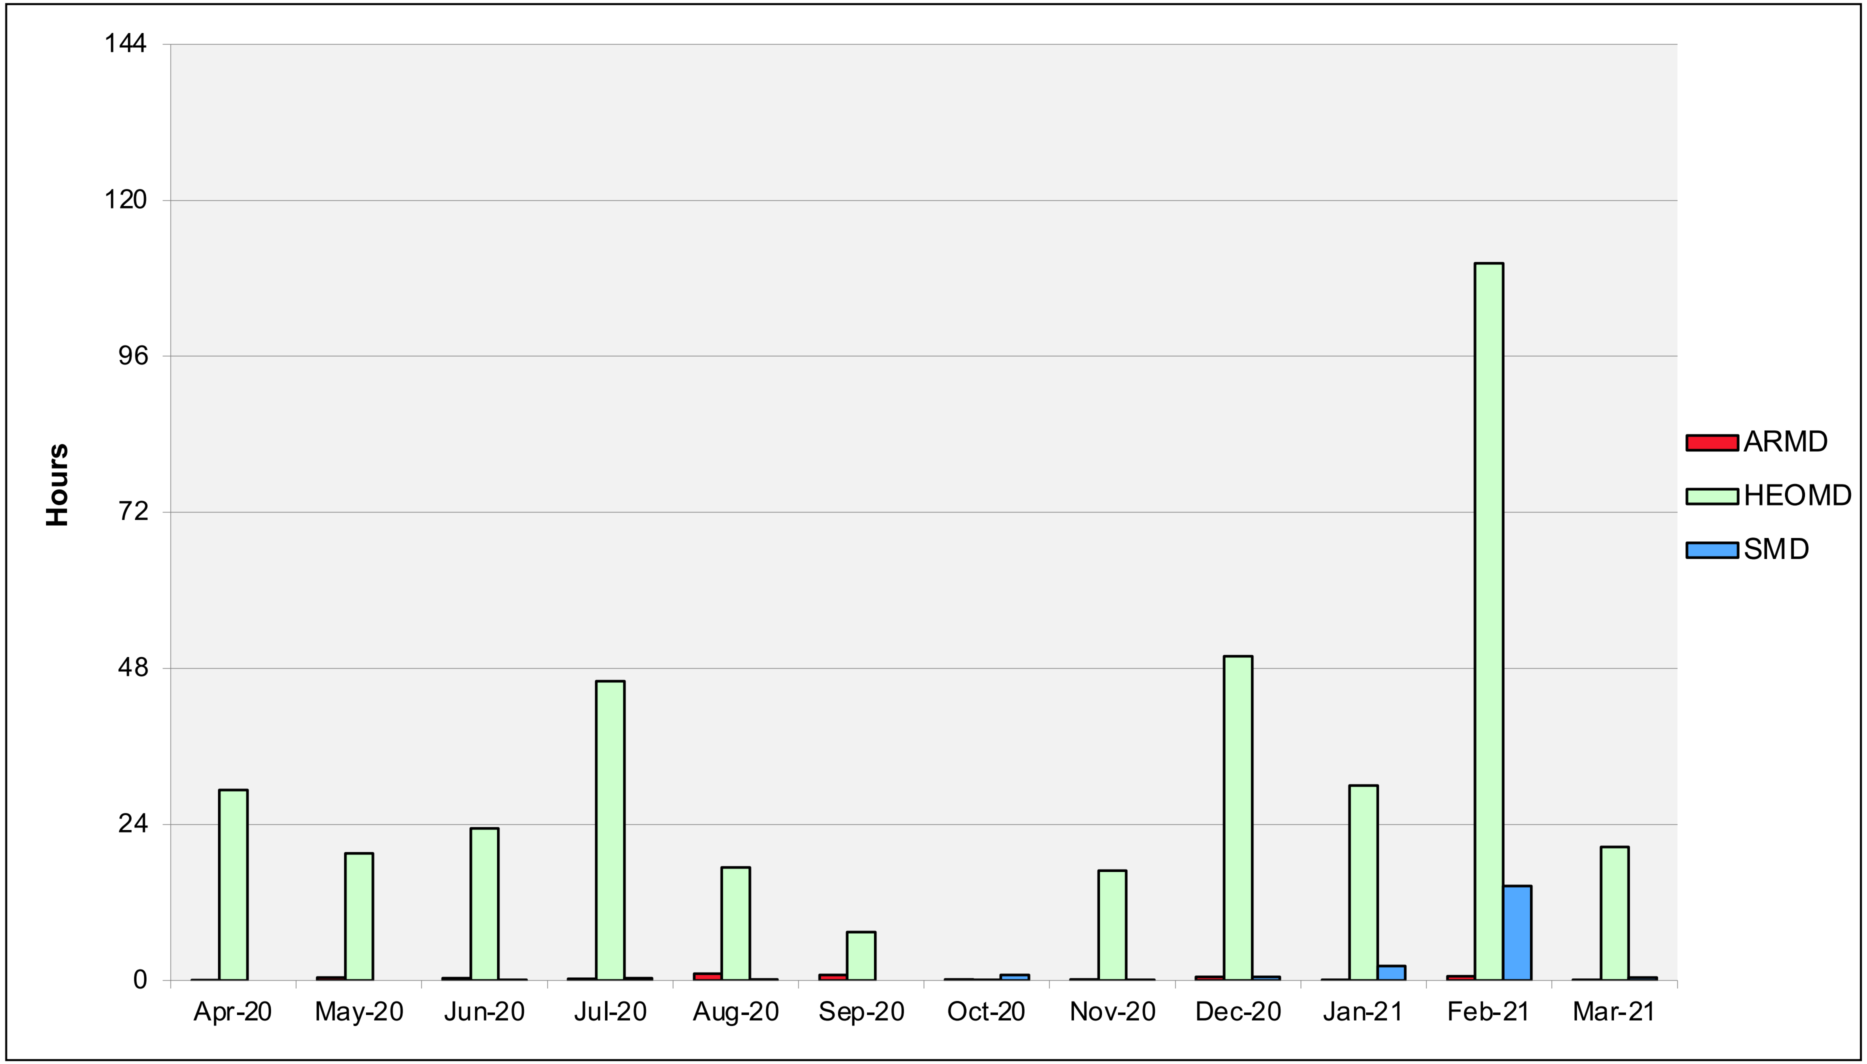

Endeavour: Average Expansion Factor